OnCommand Unified Manager 7.0

Online Help

July 2016 | 215-11255_A0
doccomments@netapp.com

NetApp®

Contents

Introduction to OnCommand Unified Manager

OnCommand Unified Manager is a suite of applications that enables you to manage your storage systems from a single interface. You can access Unified Manager and Performance Manager from the unified interface. You do not need to log in every time you navigate to the other application.

Unified Manager 7.0 and Performance Manager 7.0 share a unified management interface with the following features:

- Provides a single URL to manage both products.
- Displays cluster health and performance attributes through a central dashboard.
- Provides a single location to configure the clusters, user accounts, and authentication attributes that apply across both products.
- Eliminates multiple logins when switching from one product to the other.
- Provides a Favorites dashboard that enables easy access to the storage objects that you access frequently.

Note: Unified Manager and Performance Manager are installed individually on separate servers.

You can continue to install Unified Manager in a standalone configuration where it is not paired with Performance Manager.

OnCommand Unified Manager features

Unified Manager is built on a server infrastructure that delivers scalability, supportability, and enhanced monitoring and notification capabilities. Unified Manager supports monitoring of ONTAP 9.0, 8.3.2, 8.3.1, 8.3.0, and 8.2.x systems.

Unified Manager includes the following features:

- Discovery, monitoring, and notifications for systems that are installed with ONTAP software:
  - Physical objects: cluster nodes, disks, disk shelves, SFO pairs, ports, and Flash Cache
  - Logical objects: clusters, Storage Virtual Machines (SVMs), aggregates, volumes, LUNs, qtrees, LIFs, Snapshot copies, junction paths, NFS exports, CIFS shares, user and group quotas, and initiator groups
  - Protocols: CIFS, NFS, FC, iSCSI, and FCoE
  - Storage efficiency: SSD aggregates, Flash Pool aggregates, deduplication, and compression
  - Protection: SnapMirror relationships and SnapVault relationships
- Enhanced performance monitoring, including full integration with OnCommand Performance Manager
- Viewing the cluster discovery and monitoring status
- MetroCluster configuration: viewing and monitoring the configuration, MetroCluster switches and bridges, issues, and connectivity status of the cluster components
- Enhanced UI and comprehensive cluster visualization
- Enhanced alerts, events, and threshold infrastructure
- LDAP and local user support
- RBAC (only for a predefined set of roles)

- AutoSupport and support bundle
- Enhanced logging
- Enhanced dashboard to show capacity, availability, protection, and performance health of the environment
- Volume move interoperability, volume move history, and junction path change history
- Scope of Impact area that graphically displays the resources that are impacted for events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events
- Possible Effect area that displays the effect of the MetroCluster events
- Suggested Corrective Actions area that displays the actions that can be performed to address events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events
- Resources that Might be Impacted area that displays the resources that might be impacted for events such as for the Volume Offline event, the Volume Restricted event, and the Thin-Provisioned Volume Space At Risk event
- Support for SVMs with Infinite Volume:
  - Monitoring Infinite Volumes
  - Configuring Infinite Volume and storage class thresholds
  - Configuring rules and data policy
- Support for monitoring node root volumes
- Enhanced Snapshot copy monitoring, including computing reclaimable space and deleting Snapshot copies
- Annotations for storage objects
- Report creation and management of storage object information such as capacity, utilization, and related events
- Integration with OnCommand Workflow Automation to execute workflows for storage classes and to monitor SVMs with Infinite Volume that do not have storage classes
- The Storage Automation Store contains NetApp-certified automated storage workflow packs developed for use with OnCommand Workflow Automation (WFA). You can download the packs, and then import them to WFA to execute them. The automated workflows are available at the following link: *Storage Automation Store*

Related concepts

What you can do with Unified Manager on page 16

What you can do with Unified Manager

OnCommand Unified Manager helps you to monitor a large number of systems running clustered Data ONTAP 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, monitoring performance, configuring and managing of Infinite Volumes, annotating storage objects, 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.

Related concepts

What the Unified Manager server does on page 33

Clusters dashboard

The Clusters dashboard is the central dashboard of the integrated interface of OnCommand Unified Manager, which provides an overview of all the clusters in your storage management system. You can also navigate from the Clusters dashboard to a cluster's health dashboard or performance dashboard.

The Clusters dashboard consists of two major sections: Managed Clusters (on the left) and Cluster Details (on the right). The ⚙ command button enables you to edit, rediscover, or delete a cluster from Unified Manager.

Managed Clusters section

Lists all the clusters in your storage management system. The following details are provided for each cluster in the list:

- Cluster status icon: The status can be Critical (✖), Error (!), Warning (⚠), or Normal (✓).
- IP address or host name: Provides the host name of the cluster, and the IP address or FQDN.
- Cluster Health: Provides information about the health of the cluster as monitored by Unified Manager.
  The health status can have one of the following values: OK, OK with suppressed, Degraded, and Components not reachable.
- Pairing Status: Provides information about the connection status between this Unified Manager server and an instance of Performance Manager.
  The pairing status can have one of the following values: Good, Bad, or Not Paired.

Pairing status	Description
Good	Unified Manager is successfully paired with an instance of Performance Manager.

Pairing status	Description
Not Paired	◦ The cluster exists only in Unified Manager and has not been added to a Performance Manager server. ◦ The cluster was added to a Performance Manager server recently, but Unified Manager has not yet received the initial heartbeat signal from Performance Manager. ◦ A cluster that had been added to Performance Manager has been removed from Performance Manager. ◦ Performance Manager is down, or is in the process of restarting. ◦ The “ocie” service is down. ◦ A Performance Manager restore task is currently in progress.
Bad	◦ Unified Manager is paired with an instance of Performance Manager, but the recurring heartbeat signal has not been received for some time. ◦ Unified Manager cannot access Performance Manager

Details section

Displays the Add Cluster dialog box when no clusters are added to Unified Manager. If you have added one or more clusters, when you select a cluster, the Details section displays the Cluster: <cluster_name> dialog box.

The Details section provides information about the monitoring status, capacity, and performance of the selected cluster.

- Monitoring Status: Displays the ongoing health (Unified Manager) and performance (Performance Manager) monitoring status.
  The monitoring status can have the following values: Discovering, Poll completed, Poll failed, or Not available. The monitoring status displays an error message when the corresponding monitoring job (health or performance) fails.
- Capacity: Displays the total, used, and free aggregate capacity of the selected cluster.
- Performance: Displays the average operating speed of the cluster in number of IOPS (input/output operations per second), and the average throughput of the selected cluster in megabytes per second.

The Details section also provides navigation links to the individual cluster details pages of the OnCommand Unified Manager applications:

- The Health link navigates to the health details page of the selected cluster in Unified Manager.
- The Performance link navigates to the performance details page of the selected cluster in Performance Manager.
  If the selected cluster is not associated with an instance of Performance Manager, the Performance (Setup) option is displayed. You can click this link to associate this cluster with an instance of Performance Manager.

There may be cases where the Health or Performance link is not clickable, and the (Setup) option is unavailable. Typically this means that the cluster has been added, but that either Unified Manager or Performance Manager has not completed its initial discovery of the cluster.

Health dashboard

The Health dashboard provides an overview of the overall health status of the capacity, availability, performance, and protection health of your storage system. This dashboard also provides information on any specific issues about the storage objects.

Understanding the dashboard on page 45

Performance dashboard

The Performance dashboard provides high-level overview of the performance status for all the clusters that are being monitored by OnCommand Performance Manager. Clusters that have performance issues are displayed at the top of the Cluster list in order of severity.

The information on the Performance dashboard is updated automatically at each performance statistics collection interval. See the Performance Manager online help and documentation for more information.

Favorites dashboard

The Favorites dashboard provides a view of all the storage objects that you have marked as favorites from the storage object detail's page of OnCommand Unified Manager. You can mark storage objects that require frequent monitoring as favorites in order to avoid going through all the storage objects in your storage system each time you want to make any configuration changes.

The Favorites dashboard initially displays the steps for marking storage objects as favorites in OnCommand Unified Manager.

After you mark storage objects in your storage management system as favorites, the Favorites dashboard displays the favorite cards in chronological order. The most recent favorite card is displayed at the top left corner of the My Favorite Storage objects page. The favorite cards are refreshed every 3 minutes so that you get the latest information about your favorite objects.

The favorite cards provide the following information:

- Details about the capacity, protection, and performance of individual storage objects.
- Object hierarchy details when you position your cursor over each card.
- Links to capacity and protection information of Unified Manager or Performance Manager.

Managed Clusters page

The Managed Clusters page enables you to configure clusters in OnCommand Unified Manager. This page provides information about all the clusters in your storage management system and enables you to modify the cluster settings.

Tasks you can perform from this page

- Add, edit, or remove a cluster.
- Rediscover cluster to gather the latest data of a selected cluster.

Adding a cluster

You can add a cluster to OnCommand Unified Manager to obtain information such as the health, capacity, and configuration of the cluster so that you can find and resolve any issues that might occur. *Adding clusters* on page 20

You can also change cluster settings such as the IP address or the Performance Manager instance associated with the cluster. *Editing clusters* on page 22

Rediscovering a cluster

You can manually rediscover a cluster to obtain the latest information about the health, monitoring status, and pairing status of the cluster. You can also 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. You can access the manual rediscover option for each cluster from the ⚙ menu.

Related tasks

Adding a cluster to a Performance Manager server on page 22
Editing clusters on page 22
Moving a cluster from one Performance Manager server to another on page 24
Rediscovering clusters on page 26

Related references

Clusters dashboard on page 17

Adding clusters

You can add a cluster to OnCommand Unified Manager to obtain cluster information such as the health, capacity, and configuration of the cluster so that you can find and resolve any issues that might occur. You can also view the cluster discovery status and monitor the performance of the cluster if you associate an Performance Manager instance with the cluster.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have the following information:
    - Host name or cluster-management IP address

    The host name is the FQDN or short name that Unified Manager uses to connect to the cluster. The host name must resolve to the cluster-management IP address.
    The cluster-management IP address must be the cluster-management LIF of the administrative Storage Virtual Machine (SVM). If you use a node-management LIF, the operation fails.

  - Data ONTAP administrator user name and password
    This account must have the *admin* role with Application access set to *ontapi*, *ssh*, and *http*.
  - Type of protocol (HTTP or HTTPS) that can be configured on the cluster and the port number of the cluster

- The Unified Manager FQDN must be able to ping Data ONTAP.
  You can verify this by using the following Data ONTAP command: `ping -node node_name -destination Unified_Manager_FQDN`.

About this task

For a MetroCluster configuration, you must add both the local and remote clusters, and the clusters must be configured correctly.

Steps

1. Click ▦▾ > **Dashboard**.
2. From the **Managed Clusters** page, click **Add Cluster**.
3. In the **Add Cluster** page, specify the required values, such as the host name or IP address (IPv4 or IPv6) of the cluster, user name, password, protocol for communication, and port number.

   By default, the HTTPS protocol is selected.

   You can change the cluster-management IP address from IPv6 to IPv4 or from IPv4 to IPv6. The new IP address is reflected in the cluster grid and the cluster configuration page after the next monitoring cycle is complete.
4. In the **Link Performance Manager** section, select the name of the Performance Manager instance to which you want the cluster to be assigned.

   You can associate an instance of Performance Manager either while adding a cluster or while modifying the cluster configuration.
5. Click **Save**.
6. If HTTPS is selected, perform the following steps:
   a. In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information about the cluster.
   b. Click **Yes**.

      Unified Manager checks the certificate only when the cluster is added initially. Unified Manager does not check the certificate for each API call to Data ONTAP.

      If the certificate has expired, you cannot add a new cluster. You must first renew the SSL certificate and then add the cluster.

      The cluster is added to the Unified Manager database after the default monitoring interval of approximately 15 minutes. If the cluster is associated with an instance of Performance Manager, the cluster is automatically added to Performance Manager.

Related tasks

Adding a user on page 379
Removing a cluster in case of discovery failure on page 146

Restarting cluster discovery on page 146

Related information

NetApp KB Article 1014389: How to renew an SSL certificate in clustered Data ONTAP

Editing clusters

You can use Unified Manager to modify the settings of an existing cluster, such as the host name or IP address (IPv4 or IPv6), user name, password, protocol (HTTP to HTTPS), and port. For example, you can change the protocol from HTTP to HTTPS using the Edit Cluster dialog box.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can change the cluster management IP address from IPv6 to IPv4 or from IPv4 to IPv6. The new IP address is reflected in the cluster grid and the cluster configuration page after the next monitoring cycle finishes.

Steps

1. On the **Clusters** page, click ⚙ > **Edit** for the cluster that you want to edit.
2. In the **Edit Cluster** dialog box, modify the values of host name, IP address, protocol, port, or user name as required.

   You can also add an instance of Performance Manager to the selected cluster or change the instance of Performance Manager associated with the selected cluster.
3. Click **Save**.
4. If you have selected the HTTPS protocol, perform the following steps:
   a. In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information of the cluster.
   b. Click **Yes**.

Related tasks

Adding a user on page 379

Adding a cluster to a Performance Manager server

You add a cluster to a Performance Manager server so that you can monitor cluster performance. You can add the cluster to Performance Manager at the same time as you add it to Unified Manager using the Add Cluster page, or you can add the cluster to Performance Manager later, using the Edit Cluster page.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The Performance Manager server where you want to add the cluster must be installed with Performance Manager version 2.1 software, or later, and running.

- During the installation of Performance Manager software, you specified that you would connect the Performance Manager server with a specific Unified Manager server, and you are currently logged in to this Unified Manager server.
- The user name and password used to access the cluster must have the *admin* role with Application access set to *ontapi*, *ssh*, and *http*.

About this task

A cluster should be managed by only one instance of Performance Manager.

The process of adding a cluster to Performance Manager varies depending on whether you are adding the first cluster to the Performance Manager server or clusters already exist on that server. When adding the first cluster, you must perform Performance Manager initialization tasks. Both procedures are described in the following steps.

A single instance of Performance Manager supports a specific number of clusters and storage objects. If Performance Manager is monitoring an environment that exceeds the supported configuration, it might have difficulty collecting and analyzing configuration and performance data from the clusters. See the *OnCommand Performance Manager Release Notes* for the number of clusters, nodes, and volumes that Performance Manager can reliably support.

Steps

1. Use a web browser to log in to the Unified Manager web UI, using the IP address or URL and an appropriate user name and password.
2. From the **Managed Clusters** list, select the cluster you want to add, and then click ⚙ **> Edit**.

   The Edit Cluster page is displayed in the right pane.
3. From the **Link Performance Manager** section, select the Performance Manager server that will monitor the cluster.
4. Click **Save**.

   Note: If you receive an error message that the cluster add operation failed because of an HTTPS certificate error, ensure that you rebooted the Performance Manager server after pairing Performance Manager with Unified Manager.
5. If you selected the HTTPS protocol, perform the following steps:
   a. In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information of the cluster.
   b. Click **Yes** to authorize Performance Manager to communicate with the cluster.

   The result depends on whether the Performance Manager server is initialized:

   - If the server is already initialized, the cluster is added to the server.
     After the initial cluster inventory and data collection has completed, which might take up to 30 minutes, performance statistics are displayed in the UI.
   - If the server is not initialized, a new browser window is displayed.
6. Follow the instructions in the new browser window to set up email and AutoSupport:
   a. Specify an initial email recipient to which email alerts will be sent, and the SMTP server that will handle email communications.
   b. Specify whether AutoSupport is enabled to send information about your Performance Manager installation to technical support.

7. Click **Save and Complete Initialization**.
8. Return to the **Edit Cluster** page in the original browser window.
9. Click **Save**.

Result

After all of the objects are discovered, Performance Manager gathers historical performance data for the previous 24 hours. This enables you to view a full day of historical performance information for a cluster immediately after it is added. After the historical data is collected, real-time cluster performance data is collected, by default, every five minutes.

Moving a cluster from one Performance Manager server to another

The number of clusters that a single instance of Performance Manager can support depends on the number on the number of nodes, volumes, and other storage objects within each cluster. When too many clusters and storage objects are being monitored by a single instance of Performance Manager, you might need to move some clusters to a different instance of Performance Manager.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The Performance Manager server where you want to add the cluster is installed with Performance Manager version 2.1 software, or later, and it is running.
- During the installation of Performance Manager software, you specified that you would connect the Performance Manager server with a specific Unified Manager server, and you are currently logged in to this Unified Manager server.

About this task

Moving a cluster consists of removing the cluster from one instance of Performance Manager and adding the cluster to another instance of Performance Manager.

> **Attention:** This task is disruptive, because all cluster performance data, including historical data, storage services, and all associated events, is deleted from the original Performance Manager server.

The process of adding a cluster to the new Performance Manager server varies depending on whether you are adding the first cluster to the Performance Manager server or if one cluster has already been added to the server. When adding the first cluster, you must perform Performance Manager initialization tasks.

Steps

1. From the **Managed Clusters** list, select the cluster you want to move and click ⚙ **> Edit**.

   The Edit Cluster page displays in the right pane.
2. From the **Link Performance Manager** section, select the option **None** from the **Select Application Instance** list.
3. Click **Save**.

4. Click **Yes** in response to the warning message that all performance data for this cluster will be lost after removing the cluster.
5. From the **Link Performance Manager** section, select the new Performance Manager server that will monitor the cluster.
6. Click **Save**.
7. If you have selected the HTTPS protocol, perform the following steps:
    a. In the **Authorize Host** dialog box, click **View Certificate**.
    b. Click **Yes** to authorize Performance Manager to communicate with the cluster.

    Depending on whether the Performance Manager server is initialized, the following actions occur:

    - If the server is already initialized, the cluster is added to the server.
      After the initial cluster inventory and data collection is complete, which might take up to 30 minutes, performance statistics display in the UI.
    - If the server is not initialized, a new browser window displays.
8. Follow the instructions in the new browser window to set up email and AutoSupport:
    a. Specify an initial email recipient to which email alerts will be sent, and the SMTP server that will perform email communications.
    b. Specify whether AutoSupport is enabled to send information about your Performance Manager installation to technical support.
9. Click **Save and Complete Initialization**.
10. Return to the **Edit Cluster** page in the original browser window.
11. Click **Save**.

    The cluster is added to the server. After the initial cluster inventory and data collection is complete, performance statistics display in the UI.

Removing clusters

You can remove a cluster from OnCommand Unified Manager from the Managed Clusters page when you want to decommission a storage system or there is critical failure for the cluster.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

All the data collected for the cluster from Unified Manager and Performance Manager is lost when a cluster is deleted.

Steps

1. On the **Managed Clusters** page, click ⚙ for the cluster you want to delete.
2. Click **Remove**.
3. Click **Yes** to confirm the delete request.

The cluster is removed from Unified Manager as well as the associated Performance Manager instance.

Related tasks

Adding a user on page 379

Rediscovering clusters

You can manually rediscover a cluster from the Managed Clusters page in order to obtain the latest information about the health, monitoring status, and pairing 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.

About this task

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

Steps

1. From the **Managed Clusters** page, click ⚙ for the cluster that you want to rediscover.
2. Click **Rediscover**.

   Unified Manager rediscovers the selected cluster and displays the latest health and pairing status of the cluster in the Managed Clusters page.

   Note: You can obtain the monitoring status of the cluster from the right pane of the Clusters dashboard.

Understanding the workspace

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 OnCommand 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

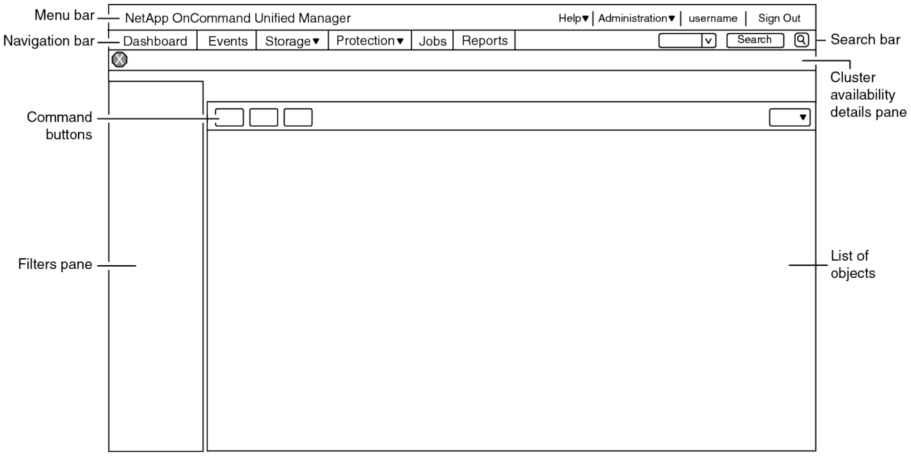

Details window layout

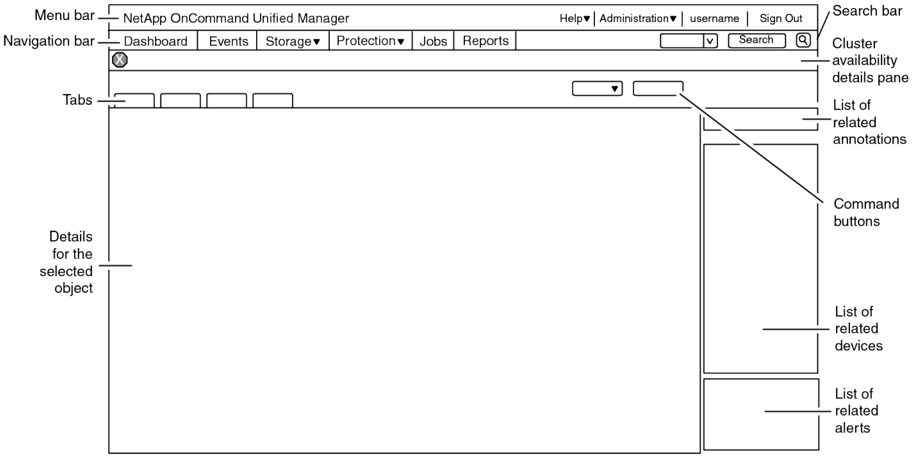

Related references

Window layout customization on page 28

Window layout customization

OnCommand Unified Manager enables you to customize the layout of information on the storage object pages. By customizing the windows, you can control which data is viewable or how it 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 apply filters to customize the display of information on the storage object pages so that only those entries that match the conditions that are provided are displayed. You can apply filters either from the Filters pane or on the columns.

The Filters pane enables you to filter some of the columns based on the options that are selected. For example, on the Volumes page, you can use the Filters pane to filter only the Status and State columns. To display all volumes that are offline, you can select the appropriate filter option under State.

Alternatively, you can choose to set filters on the columns by using the filter icon ( ). You can then use the wildcard character filter (?) or wildcard string filter (*) to narrow your search. For example, on the Volumes page, you can search for a volume, vol234, by using the string filter in the Volume column. You can type `*vol`, and all the volumes with names containing “vol” are listed. You can type `vol?` to view the list of all volumes with the name containing “vol” followed by one more character—for example, vol1 or vol2. You can type `vol` to view the list of all the volumes with names that start with “vol”.

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 Aggregates page 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 Aggregates page 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%. For columns that display capacity data in percentage, you can view values up to four decimal points by moving your mouse pointer over the value that is displayed in the column.

Hiding or redisplaying the columns

You can click the column display icon ( ) to select which columns you want to display.

Exporting data

You can click the export icon ( ) to export data to a comma-separated values (`.csv`) file and use the exported data to build reports.

How graphs of performance data work

Performance Manager uses graphs or charts to show you volume performance statistics and events over a specified period of time.

The graphs enable you to customize the range of time for which to view data. The data is displayed with the time frame on the horizontal axis of the graph and the counters on the vertical axis, with point intervals along the graph lines. The vertical axis is dynamic; the values adjust based on the peaks of the expected or actual values.

Selecting time frames

On the Volume Details page, the Historic data chart enables you to select a time frame for all graphs on the page. The 1d, 5d, 10d, and 30d buttons specify 1 day through 30 days (1 month) and the **Custom** button enables you to specify a custom time range within that 30 days. Each point on a graph represents a 5-minute collection interval, and a maximum of 30 days of historical performance data is retained. Note that intervals also account for network delays and other anomalies.

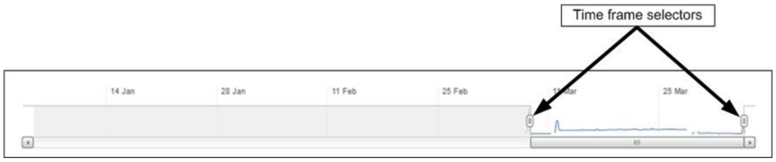

In this example, the Historic data chart has a time frame set to the beginning and the end of the month of March. In the selected time frame, all historic data before March is grayed out.

Viewing data point information

To view data point information on a graph, you can position the mouse cursor over a specific point within the graph, and a pop-up box displays listing the value and date and time information.

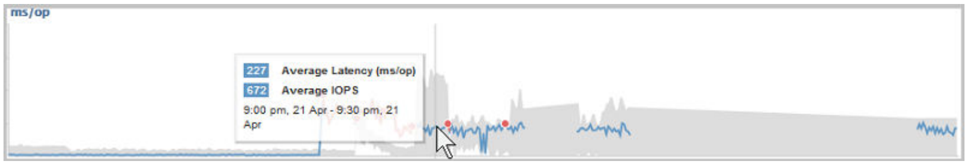
ms/op
227 Average Latency (ms/op)
672 Average IOPS
9:00 pm, 21 Apr - 9:30 pm, 21 Apr

In this example, positioning the mouse cursor over the IOPS chart on the Volume Details page displays the response time and operations values between 3:50 a.m. and 3:55 a.m. on October 20th.

Viewing event information

To view event information on a graph, you can position your mouse pointer over an event icon to view summary information in a pop-up box, or you can click the event icon for more detailed information.

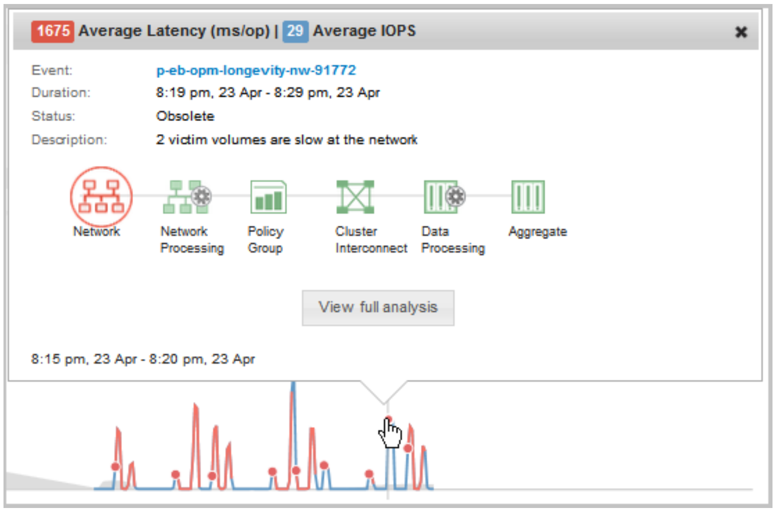

1675 Average Latency (ms/op) | 29 Average IOPS

Event: p-eb-opm-longevity-nw-91772
Duration: 8:19 pm, 23 Apr - 8:29 pm, 23 Apr
Status: Obsolete
Description: 2 victim volumes are slow at the network

8:15 pm, 23 Apr - 8:20 pm, 23 Apr

In this example, on the Volume Details page, clicking an event icon on the Latency chart displays detailed information about the event in a pop-up box. The event is also highlighted in the Events List.

Using the Unified Manager Help

The Help includes information about all features included in OnCommand 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.

About this task

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

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

Choices

- To learn about specific fields or parameters, click ?.
- To view all the Help contents, click **Help** > **Help Contents**.

  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.

Related tasks

Bookmarking your favorite Help topics on page 31

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**.

Exporting data to CSV files

You can export data to a comma-separated values (`.csv`) file, and 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 Events page to create a report, and then take appropriate action.

Before you begin

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

About this task

You can export data to a `.csv` file from the following pages:

- Events page
- Clusters page
- Nodes page
- Storage Virtual Machines page
- Aggregates page
- Volumes page

The export functionality is not supported for the constituents of an Infinite Volume—you cannot export details of the constituents to a `.csv` file.

Steps

1. Perform one of the following actions:

If you want to...	Do this...
Export event details	Click **Events**.
Export storage object details	Click **Storage**, and then select an object. You can select clusters, nodes, SVMs, aggregates, or volumes.

2. Click **Export**.
3. Click **Yes** to confirm the export request.
4. In the dialog box that is displayed, select the appropriate application to open the `.csv` file.

   It might take a while for the data to be exported.

5. Click **OK**.

Related references

Volumes page on page 211
Aggregates page on page 200
Storage Virtual Machines page on page 184
Clusters page on page 168

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.

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 OnCommand Unified Manager:

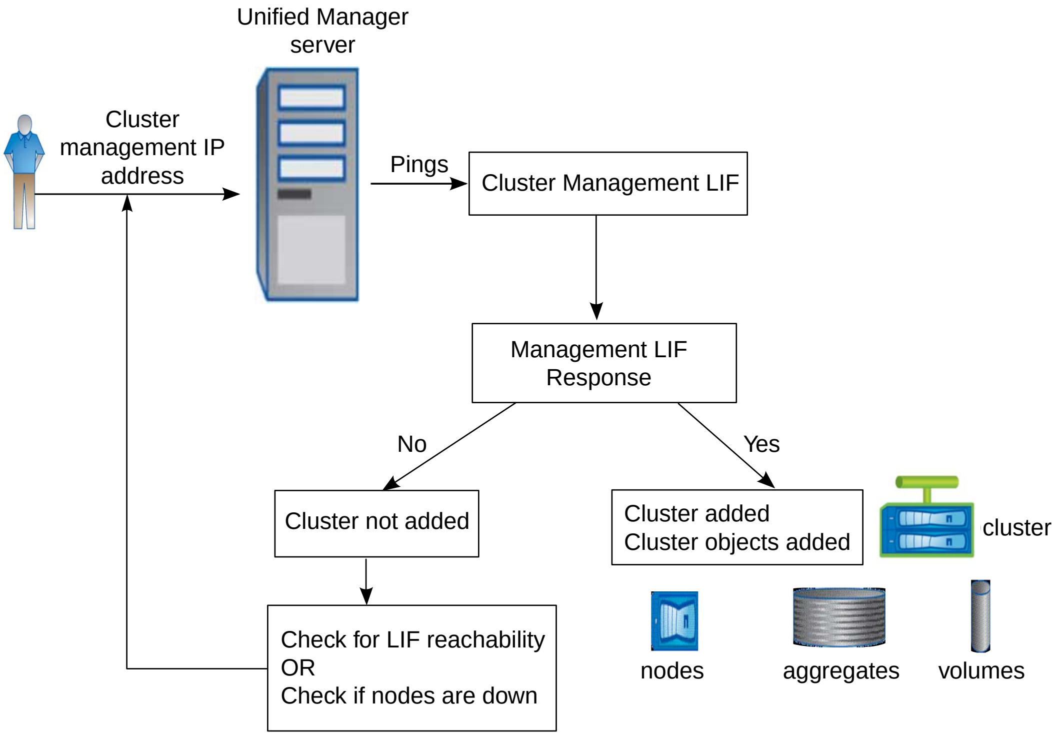

Searching for storage objects

You can use the search bar to find your storage objects. Search results are sorted by storage object type, and you can filter them using the drop-down menu. A valid search must contain at least three characters.

Before you begin

You must have one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.

Step

1. Type your search parameters into the search bar and press **Enter**.

   You can use the filter to select a specific storage object type for your search.

   #### Example

   If you want to search for one of your aggregates, select **Aggregates** from the filter, and then type the name of the aggregate or type any three characters in the aggregate's name in the search bar. You can then select the appropriate aggregate from the drop-down list.

How timestamps work in Unified Manager

The timestamp displayed in Unified Manager is based on your web browser's time zone and not on the time zone that is configured for the Unified Manager server.

For example, the time that is displayed in the dashboard when an event is generated or the estimated end time for a volume move operation is based on the browser time zone on which you have launched Unified Manager web UI.

Getting started

After you deploy the Unified Manager virtual appliance, you have to perform several configuration tasks in the setup wizard before you start monitoring your clusters. The initial configuration tasks include setting the host name, setting up alert notifications, adding users, setting up email and time zone, and enabling AutoSupport.

Configuring your environment after deployment

After you deploy and install Unified Manager, there are several configuration tasks that you might want to perform before you start monitoring your clusters, such as changing the host name, adding alerts, and adding users.

Before you begin

- You must have installed Unified Manager, and completed the Unified Manager initial setup.
- You must have the OnCommand Administrator role.

About this task

After you complete the Unified Manager initial setup, you can add clusters. If you did not add clusters after the initial setup, you must add clusters before you can start monitoring cluster objects. You can add clusters at any time. However, there are some configuration changes that you might want to make to Unified Manager before or after adding clusters.

Choices

- *Changing the Unified Manager host name* on page 35

  When you deployed Unified Manager, an SSL certificate was generated for HTTPS access. A host name was associated with the certificate, allowing you to use the host name to access the Unified Manager web UI. You might want to change this host name after deployment.

- *Configuring Unified Manager to send alert notifications* on page 38

  After clusters are added to Unified Manager, you can monitor them, but you cannot receive notifications about events in your cluster environment until you configure several options (for example, the email address from which notifications are sent, and the users who should receive the alerts). You might also want to modify the default threshold settings at which events are generated.

Changing the Unified Manager host name if Unified Manager is installed as a virtual appliance

The network host is assigned a name when the virtual appliance is first deployed (if Unified Manager is installed as a virtual appliance). You can change the host name after deployment. If you change the host name, you must also regenerate the HTTPS certificate.

Before you begin

You must be logged in to Unified Manager as the maintenance user, or have the OnCommand Administrator role assigned to you to perform these tasks.

About this task

> **Note:** If Unified Manager is installed on Red Hat Enterprise Linux, you should not perform this workflow to change the host name. Instead, you must change the host name through Red Hat Enterprise Linux commands.

You can use the host name (or the host IP address) to access the Unified Manager web UI. If you configured a static IP address for your network during deployment, then you would have designated a name for the network host. If you configured the network using DHCP, the host name should be taken from the DNS. If DHCP or DNS is not properly configured, the host name “OnCommand” is automatically assigned and associated with the security certificate.

Regardless of how the host name was assigned, if you change the host name, and intend to use the new host name to access the Unified Manager web UI, you must generate a new security certificate.

If you access the web UI by using the server's IP address instead of the host name, you do not have to generate a new certificate if you change the host name. However, it is the best practice to update the certificate so that the host name in the certificate matches the actual host name.

If you change the host name in Unified Manager, you must manually update the host name in OnCommand Workflow Automation (WFA). The host name is not updated automatically in WFA.

The new certificate does not take effect until the Unified Manager virtual machine is restarted.

Steps

1. *Generate an HTTPS security certificate* on page 406

   If you want to use the new host name to access the Unified Manager web UI, you must regenerate the HTTPS certificate to associate it with the new host name.

2. *View the HTTPS security certificate* on page 405

   You must verify that the correct information is displayed after generating a new security certificate.

3. *Restart the Unified Manager virtual machine* on page 405

   After you regenerate the HTTPS certificate, you must restart the Unified Manager virtual machine.

Changing the OnCommand Unified Manager host name in Red Hat Enterprise Linux

At some point, you might want to change the host name of the Red Hat Enterprise Linux machine on which you have installed Unified Manager. For example, you might want to rename the host to more easily identify your Unified Manager servers by type, workgroup, or monitored cluster group when you list your Red Hat Enterprise Linux machines.

Before you begin

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

About this task

You can use the host name (or the host IP address) to access the Unified Manager web UI. If you configured a static IP address for your network during deployment, then you would have designated a name for the network host. If you configured the network using DHCP, the host name should be taken from the DNS server.

Regardless of how the host name was assigned, if you change the host name and intend to use the new host name to access the Unified Manager web UI, you must generate a new security certificate. If you access the web UI by using the server's IP address instead of the host name, you do not have to generate a new certificate if you change the host name. However, it is the best practice to update the certificate, so that the host name in the certificate matches the actual host name. If you change the host name in Unified Manager, you must manually update the host name in Workflow Automation. The host name is not updated automatically. The new certificate does not take effect until the Red Hat Enterprise Linux machine is restarted.

> **Important:** If connections that enable performance monitoring are currently configured between the Unified Manager server and one or more Performance Manager servers, executing this task invalidates those connections and deactivates any further performance monitoring updates from Performance Manager servers to the Unified Manager UI. You must reactivate those connections after completing this task.

Steps

1. Log in as the root user to the Unified Manager Red Hat Enterprise Linux machine that you want to modify.
2. Stop the Unified Manager software and the associated MySQL software by entering the following commands in the order shown:

   ```
   service ocieau stop
   ```

   ```
   service ocie stop
   ```

   ```
   service rp stop
   ```

   ```
   service mysqld stop
   ```

3. Edit the HOSTNAME parameter in the `/etc/sysconfig/network` file to specify the new fully qualified domain name (FQDN), and save the file:

   ```
   HOSTNAME=new_FQDN
   ```

   Example

   ```
   HOSTNAME=nuhost.corp.widget.com
   ```

4. If there is an entry in the `/etc/hosts` file listing your IP address with the old host name, change it to the new name:

   ```
   ip-address new_FQDN new_hostname
   ```

   Example

   ```
   10.10.10.54 nuhost.corp.widget.com nuhost
   ```

5. Change the host name using the Linux `hostname` command:

   ```
   hostname new_FQDN
   ```

   Example

   ```
   hostname nuhost.corp.widget.com
   ```

6. Regenerate the HTTPS certificate for the server:

   ```
   /opt/netapp/essentials/bin/cert.sh create
   ```

7. Restart the network service:

   ```
   service network restart
   ```

8. After the service is restarted, verify whether the new host name is able to ping itself:

   ```
   ping new_hostname
   ```

Example

```
ping nuhost
```

This command should return the same IP address that was set earlier for the original host name.

9. After you complete and verify your host name change, restart Unified Manager by entering the following commands in the order shown:

```
service mysqld start
service rp start
service ocie start
service ocieau start
```

Configuring Unified Manager to send alert notifications

You can configure Unified Manager to send notifications that alert you about events in your environment. Before notifications can be sent, you must configure several other Unified Manager options.

Before you begin

You must have the OnCommand Administrator role.

About this task

After deploying Unified Manager and completing the initial configuration, you should consider configuring your environment to trigger alerts and generate notification emails or SNMP traps.

Steps

1. *Configure notification settings* on page 75

   If you want alert notifications sent when certain events occur in your environment, you must supply an email address from which the alert notification can be sent. If your configuration uses an SMTP server for email authentication, then you must provide the user name and password for the server. If you want to use SNMP traps, you can select that option and provide the necessary information.

2. *Enable remote authentication* on page 395

   If you want remote LDAP or Active Directory users to access the Unified Manager instance and receive alert notifications, then you must enable remote authentication.

3. *Add authentication servers* on page 398

   If you enable remote authentication, then you must identify authentication servers.

4. *Edit global threshold settings* on page 128

   You can modify the threshold settings for aggregates, volumes, and certain types of protection relationships. These settings determine when an event should be generated, which can affect when an alert notification is sent.

5. *Add users* on page 379

   You can add several different types of local or remote users and assign specific roles. When you create an alert, you assign a user to receive the alert notifications.

6. *Add alerts* on page 102

   After you have added the email address for sending notifications, added users to receive the notifications, configured your network settings, and configured SMTP and SNMP options needed for your environment, then you can assign alerts.

Customizing your environment

After you deploy the Unified Manager virtual appliance and access the web UI, you can customize the configuration of several options to meet the needs of your cluster environment.

Enabling periodic AutoSupport

You can choose to have specific, predefined messages sent automatically to technical support to ensure correct operation of your environment, and to assist you in maintaining the integrity of your environment.

Before you begin

You must be logged in as the maintenance user.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server** > **AutoSupport**.
4. To read about what periodic AutoSupport entails, click **View AutoSupport Description**.

   The dialog box also displays the product serial number, which is the number that technical support uses to find the AutoSupport messages.
5. Select the **Enable Periodic AutoSupport** check box, and then click **Save**.

Related tasks

Adding a user on page 379
Sending on-demand AutoSupport messages on page 461

Editing the global setup options

You can configure many global options that control how Unified Manager operates. This includes options for thresholds, quotas, authentication, and more. When you configure the options globally, the default values of the objects are modified. If the default values of an object have been changed at the object level, the global options of the object are not modified.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

- You can modify the following options from the Setup Options dialog box:
  - Thresholds
    You can configure the global threshold values for aggregates, volumes, and relationships.
  - Management Server
    You can configure the settings related to AutoSupport, HTTPS, and authentication.
  - General Settings
    You can configure the settings related to events and notifications.

  - Quota Settings
    You can configure the settings related to user and user group quota email notifications.
  - Add-ons
    You can configure the settings for OnCommand Workflow Automation.
- You can also modify the threshold settings for each object from the details page for that object.
- Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

Steps

1. Click ▦ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, select the category from the left navigation pane, and then select the option and modify the settings as required.
4. Click **Save and Close**.

Related tasks

Adding a user on page 379

Working with HTTPS security certificates

You can view and regenerate an existing HTTPS certificate or download and install new certificates.

Before you begin

You must be signed in to Unified Manager as the maintenance user or have the OnCommand Administrator role assigned to you to perform these tasks.

About this task

During deployment of the virtual appliance, a self-signed SSL certificate is generated and is associated with the “OnCommand” host name and a user-specified IP address. You can use this certificate, generate a new one, or download a certificate signing request and install a certificate signed by a Certificate Authority. You can also view the content of the certificate you are using.

Choices

- *Generating an HTTPS security certificate* on page 406

  You might generate a new HTTPS security certificate for multiple reasons, including if you want to sign with a different Certificate Authority or if the current security certificate has expired. The new certificate replaces the existing certificate.

- *Downloading an HTTPS certificate signing request* on page 408

  You can download a certification request for the current HTTPS security certificate so that you can provide the file to a Certificate Authority to sign. A CA-signed certificate helps prevent man-in-the-middle attacks and provides better security protection than a self-signed certificate.

- *Installing an HTTPS security certificate* on page 408

  You can upload and install a security certificate after a Certificate Authority has signed and returned it. The file that you upload and install must be a signed version of the existing self-signed certificate. A CA-signed certificate helps prevent man-in-the middle attacks and provides better security protection than a self-signed certificate.

- *Viewing the HTTPS security certificate* on page 405

  You can compare the HTTPS certificate details to the retrieved certificate in your browser to ensure that your browser's encrypted connection to Unified Manager is not being intercepted. You can also view the certificate to verify the content of a regenerated certificate or to view alternate URL names from which you can access Unified Manager.

Related references

HTTPS Setup Options dialog box on page 409

Description of Setup options dialog boxes

You can use the Setup Options dialog box to configure global threshold values for aggregates, volumes, and protection relationships; modify options related to AutoSupport, HTTPS certificates, and authentication; and to configure notification settings.

Setup Options dialog box

You can use the Setup Options dialog box to configure global threshold values, management server options, and notification settings.

You must have the OnCommand Administrator or Storage Administrator role.

- *Thresholds* on page 41
- *Management Server* on page 42
- *General Settings* on page 42
- *Quota Settings* on page 42
- *Add-ons* on page 42
- *Command buttons* on page 43

Thresholds

The Thresholds area enables you to configure the global threshold values for aggregates and volumes, and protection relationships:

- Aggregates
  Specifies the global threshold values for monitored aggregates.
  For more detailed information, see the following topic:
  Aggregate Thresholds Setup Options dialog box on page 130
- Volumes
  Specifies the global threshold values for monitored volumes.
  For more detailed information, see the following topic:
  Volume Thresholds Setup Options dialog box on page 132
- Relationships
  Specifies the global threshold values for protection relationships.
  For more detailed information, see the following topic:
  Relationships Thresholds Setup Options dialog box on page 332

Note: You should use OnCommand System Manager to set threshold values for qtrees.

Management Server

The Management Server area enables you to view or modify options related to AutoSupport, HTTPS certificates, and authentication. These are capabilities that are not directly related to storage management:

- AutoSupport
  Enables you to view the AutoSupport description, enable periodic AutoSupport, or send an on-demand AutoSupport message.
  For more detailed information, see the following topic:
  AutoSupport Setup Options dialog box on page 43
- HTTPS
  Enables you to view the current security certificate, download a certificate signing request, to generate a new HTTPS certificate, or install a new HTTPS certificate.
  For more detailed information, see the following topic:
  HTTPS Setup Options dialog box on page 409
- Authentication
  Allows you to enable authentication, configure the settings to retrieve data from authentication servers, add or delete authentication servers, and test authentication.
  For more detailed information, see the following topic:
  Authentication Setup Options dialog box on page 401

General Settings

The General Settings area enables you to configure notification settings:

- Notification
  Enables you to specify an email address, SMTP server details, and SNMP trap details in order to receive notifications. SNMPv1 and SNMPv3 are supported.
  For more detailed information, see the following topic:
  Notification Setup Options dialog box on page 107

Quota Settings

The Quota Settings area enables you to configure settings related to user and group quotas.

- Email Address Rules
  Enables you to create rules to specify the email address of a user or user group to which emails are sent when there is a breach in the quotas.
  For more detailed information, see the following topic:
  Rules to Generate User and Group Quota Email Addresses dialog box on page 141
- Email Notification Format
  Enables you to create a notification format for the emails that are sent to a user or a user group when there is a breach in the quotas.
  For more detailed information, see the following topic:
  Email Notification Format dialog box on page 140

Add-ons

The Add-ons area enables you to configure settings for OnCommand Workflow Automation.

- Set Up OnCommand Workflow Automation
  Enables you to add, modify, or delete the OnCommand Workflow Automation settings.
  For more detailed information, see the following topic:
  Set Up OnCommand Workflow Automation on page 358

Command buttons

The command buttons enable you to save or cancel the setup options:

Restore to Factory Defaults
Enables you to restore the configuration settings to the factory default values.

Save
Saves the configuration settings for the selected option.

Save and Close
Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel
Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Editing the global setup options on page 39

Related references

Rules to Generate User and Group Quota Email Address dialog box on page 141
Email Notification Format dialog box on page 140

AutoSupport Setup Options dialog box

The AutoSupport area in the Setup Options dialog box enables you to view the AutoSupport description, enable periodic AutoSupport, or send an on-demand AutoSupport message.

Information

You can perform the following operation:

View AutoSupport Description
Displays the AutoSupport description, including the customer benefits and security description.

Periodic AutoSupport

Enables you to have specific, predefined messages to technical support for issue diagnosis and resolution periodically generated.

On-Demand AutoSupport

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

Send to Technical Support
Indicates that you want to send an on-demand message to technical support for any issues that have occurred.

Send to Email Recipient
Indicates that you want to send an on-demand message to a specified recipient for any issues that have occurred.

Generate and Send AutoSupport
Enables you to generate and send an on-demand message to technical support, a specified email recipient, or both for any issues that have occurred.

Command buttons

The command buttons enable you to save or cancel the setup options:

Save

Saves the configuration settings for the selected option.

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Enabling periodic AutoSupport on page 39

Sending on-demand AutoSupport messages on page 461

Monitoring your system from the dashboard

The dashboard provides a cumulative at-a-glance information about your system. The dashboard enables you to assess the overall capacity, availability, performance, and protection health of the managed clusters, and quickly note, locate, diagnose, or assign for resolution any specific issues that might occur.

Understanding the dashboard

The Unified Manager dashboard provides cumulative at-a-glance information about your storage environment. The dashboard consists of two areas: the Quick Takes area provides information about the health of your storage objects, and the Unresolved Incidents and Risks area displays events related to the availability, capacity, performance, and protection of the storage objects.

The following image illustrates the panes that are displayed on the Unified Manager dashboard:

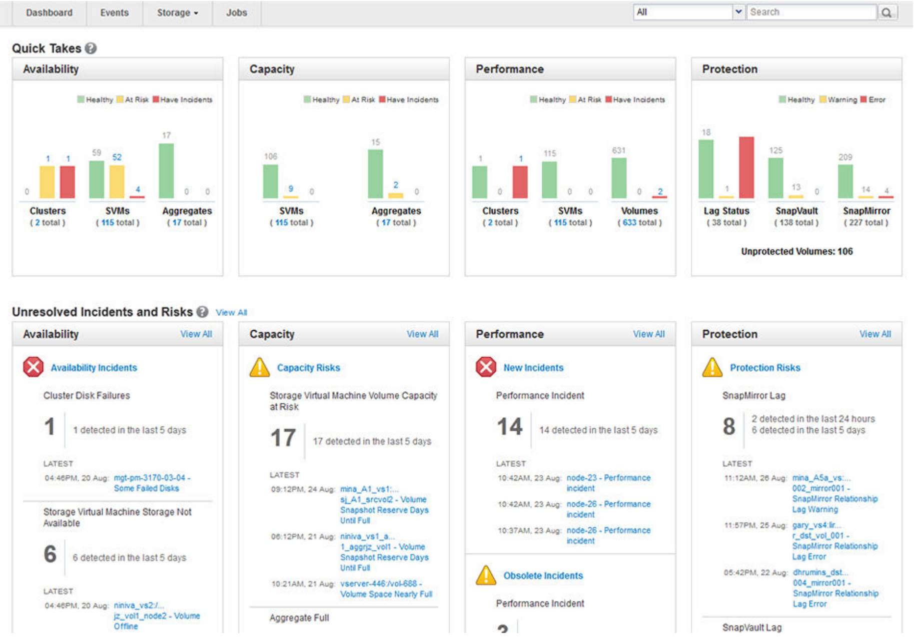

Cluster Disk Failures
1 | 1 detected in the last 5 days
LATEST
04:46PM, 20 Aug: mgt-pm-3170-03-04 - Some Failed Disks

Storage Virtual Machine Storage Not Available
6 | 6 detected in the last 5 days
LATEST
04:46PM, 20 Aug: niniva_vs2:/...jz_vol1_node2 - Volume Offline

Storage Virtual Machine Volume Capacity at Risk
17 | 17 detected in the last 5 days
LATEST
09:12PM, 24 Aug: mina_A1_vs1:...sj_A1_srcvol2 - Volume Snapshot Reserve Days Until Full
06:12PM, 21 Aug: niniva_vs1_a...1_aggrjz_vol1 - Volume Snapshot Reserve Days Until Full
10:21AM, 21 Aug: vserver-446:/vol-688 - Volume Space Nearly Full

Aggregate Full

Performance Incident
14 | 14 detected in the last 5 days
LATEST
10:42AM, 23 Aug: node-23 - Performance incident
10:42AM, 23 Aug: node-26 - Performance incident
10:37AM, 23 Aug: node-26 - Performance incident

Performance Incident

SnapMirror Lag
8 | 2 detected in the last 24 hours
6 detected in the last 5 days
LATEST
11:12AM, 26 Aug: mina_A5a_vs:...002_mirror001 - SnapMirror Relationship Lag Warning
11:57PM, 25 Aug: gary_vs4:lr...r_dst_vol_001 - SnapMirror Relationship Lag Error
05:42PM, 22 Aug: dhrumins_dst...004_mirror001 - SnapMirror Relationship Lag Error

SnapVault Lag

Quick Takes area

Displays, as a graph, information about the health of your storage objects such as clusters, aggregates, and Storage Virtual Machines (SVMs), and protection relationships. The Quick Takes area displays events generated for the following categories:

Availability

Displays information about the availability of clusters, SVMs, and aggregates that are monitored by the Unified Manager server. Based on the availability-related events that are generated, the storage objects are categorized as Healthy, At Risk, or Have Incidents.

Capacity

Displays information about the capacity of SVMs and aggregates that are monitored by the Unified Manager server. Based on the capacity-related events that are generated, the storage objects are categorized as Healthy, At Risk, or Have Incidents.

Performance

If performance monitoring connections are configured, displays information about the performance of clusters, SVMs, and volumes that are monitored by OnCommand Performance Manager and the Unified Manager server. Based on the performance-related incidents that are generated, the storage objects are categorized as Healthy or Have Incidents.

Protection

Displays information about the protection relationships that are monitored by the Unified Manager server. Based on the protection-related events that are generated, the protection relationships are categorized as Healthy, Warning, or Error.

Unresolved Incidents and Risks area

Displays new or acknowledged events related to the availability, capacity, performance, and protection of storage objects. Based on the severity, these events are categorized as incidents and risks, which help you to take corrective measures. Incidents refer to issues that have already impacted the availability, capacity, performance, or protection of storage objects. Risks refer to issues that can impact the availability, capacity, performance, or protection of storage objects.

The Unresolved Incidents and Risks area displays events generated for the following categories:

Availability

Displays information about the availability-related issues that can impact or have already impacted data availability in the storage objects.

Capacity

Displays information about the capacity-related issues that can impact or have already impacted data capacity in the storage objects.

Performance

Displays information about the performance related issues that impact or have impacted the performance health of the storage objects.

Protection

If performance monitoring connections are configured, displays information about the protection-related issues that can impact or have already impacted data protection in the storage objects. The panes also display the three most recent events that are generated.

Related references

Description of the Dashboard window on page 46

Description of the Dashboard window

You can use the Dashboard page to get a quick glance of the objects that are being monitored.

Cluster availability details pane

When a cluster is not reachable, Unified Manager displays the details in a pane at the top of every screen. If all clusters are reachable, this pane is hidden.

You can refresh the information displayed in the pane by pressing F5 or by logging out and logging back in to Unified Manager. This action ensures that the pane displays the latest information about clusters that are currently not reachable. For example, if a cluster with a Cluster Not Reachable event is removed or if the state of an event is Obsolete, information about the event is removed when you refresh the pane.

If any of the Cluster Not Reachable events are resolved, information about these events is automatically removed.

You can view detailed information about a cluster that is unreachable by clicking the **Details** button. This action opens the Events page and closes the pane automatically. After the pane is closed, it is displayed again only when you log back in to Unified Manager. You can continue to track information about clusters that are unreachable in the availability pane of the dashboard.

Quick Takes area

The Quick Takes area displays, as a graph, the health of storage objects such as clusters, aggregates, and Storage Virtual Machines (SVMs). Based on the availability, capacity, performance, and protection-related events that are generated, these storage objects are categorized as Healthy, At Risk, or Have Incidents, or as Healthy, Warning or Error for protection-related events.

Availability pane

Displays, as a graph, information about the availability of clusters, aggregates, and SVMs that are monitored by the Unified Manager server. The storage objects are categorized as Healthy, At Risk, or Have Incidents. For example, the status of a cluster that lacks spare disks is displayed as At Risk.

This pane also displays the number of storage objects in each of the categories.

Capacity pane

Displays, as a graph, information about the capacity of aggregates and SVMs that are monitored by the Unified Manager server. The storage objects are categorized as Healthy, At Risk, or Have Incidents. For example, the status of an aggregate whose used capacity has reached the full threshold value is displayed as At Risk.

This pane also displays the number of storage objects in each of the categories.

Performance pane

If Unified Manager is integrated with OnCommand Performance Manager, displays, as a graph, information about the performance of clusters, SVMs, and volumes that are monitored by OnCommand Performance Manager and the Unified Manager server. Based on the performance-related incidents that are generated, the storage objects are categorized as Healthy or Have Incidents. For example, the status of a volume whose I/O response time to its workload has reached the full threshold value is displayed as Have Incidents.

This pane also displays the total number of clusters, SVMs, and volumes that are monitored by Unified Manager server. The objects displayed are categorized as Healthy by default. Clicking on any of the object totals takes you to the page for that object. For example, clicking the cluster total takes you to the Clusters page. The total number of monitored objects is displayed regardless of whether Unified Manager is integrated with OnCommand Performance Manager.

Protection pane

Displays, as a graph, information about protection relationships that are monitored by the Unified Manager server. The protection relationships are categorized as Healthy, Warning, or Error. For example, a relationship that has a lag duration that exceeds the lag warning threshold is displayed as Warning. Protection monitoring is not supported for Data ONTAP 8.2 and earlier.

This pane also displays the total number of storage objects in each of the protection categories. Clicking the links for the lag status, SnapVault, or SnapMirror categories takes you to a filtered list of those objects in the Volume Protection Relationships page.

Related concepts

What events are on page 73

Related references

Performance event analysis and notification on page 66

Unresolved Incidents and Risks area

The Unresolved Incidents and Risks area enables you to view new and acknowledged events related to data availability, data capacity, performance, and protection status of cluster objects such as aggregates, Storage Virtual Machines (SVMs), and volumes. Based on the severity, events are categorized as incidents and risks, which helps you in taking corrective measures.

Availability pane

Displays information about availability-related issues that can impact (risk) or have already impacted (incident) the availability of data in the storage objects. For example, the area displays information about volumes that are offline or aggregates that lack spare disks.

The three most recent events are displayed by default; however, you can view all the availability events that have been generated by clicking the **View All** link.

Capacity pane

Displays information about capacity-related issues that can impact or have already impacted data capacity in the storage objects. For example, the area displays information about volumes that have reached full capacity or aggregates that are overcommitted.

The three most recent events are displayed by default; however, you can view all the capacity events that have been generated by clicking the **View All** link.

Performance pane

If performance monitoring connections are configured, displays information about the performance issues that have impacted the performance of monitored clusters, SVMs, and volumes. For example, the area displays information about volumes whose workload has reached the full threshold value.

By default, the three most recent unresolved events are displayed in the New Incidents section, and events that have been resolved or resolved on their own in the last 24 hours are displayed the Obsolete Incidents section; You can also view all the protection events that have been generated by clicking the **View All** link.

The three most recent events are displayed by default; however, you can view all the performance events that have been generated by clicking the **View All** link.

Protection pane

Displays information about protection-related issues that can impact or have already impacted data protection in the storage objects. For example, the area displays information about volume Snapshot copy reserves that are full or almost full, replicas that are out of date or nearly out of date, SnapMirror and SnapVault relationships that are broken, and protection destinations that are full or almost full.

The three most recent events are displayed by default; however, you can view all the protection events that have been generated by clicking the **View All** link.

Related concepts

What events are on page 73

Related references

Performance event analysis and notification on page 66

Managing storage objects using the Favorites option

The Favorites option enables you to manage the storage objects in OnCommand Unified Manager by marking them as favorites. You can quickly view the status of your favorite storage objects and fix issues before they become critical.

Tasks you can perform from the Favorites dashboard

- View the list of storage objects marked as favorite.
- Add storage objects to the Favorites list.
- Remove storage objects from the Favorites list.

Viewing the Favorites list

You can view the capacity, performance, and protection details of different storage objects from the Favorites list. The performance details of storage objects are displayed only if OnCommand Unified Manager is paired with OnCommand Performance Manager. The details of a maximum of 20 storage objects are displayed in the Favorites list.

Adding storage objects to the Favorites list

You can use OnCommand Unified Manager to add storage objects to the Favorites list, and then monitor these objects for health, capacity, and performance. You can only mark clusters, volumes, and aggregates as favorite.

Removing storage objects from the Favorites list

You can remove storage objects from the Favorites list in OnCommand Unified Manager when you no longer require them to be marked as favorite.

Adding storage objects to Favorites list

You can use OnCommand Management to add storage objects to a Favorites list so you can monitor the objects for health, capacity, and performance. You can use object status in the Favorites list to determine issues and fix them before they become critical. The Favorites list also provides the most recent monitoring status of a storage object.

About this task

You can add up to 20 clusters, aggregates, or volumes to the Favorites list.

Steps

1. Go to the **Details** page of the storage object that you want to mark as a favorite.
2. Click the star icon to add the storage object to the Favorites list.

Adding a cluster to the Favorites list

1. Click Clusters.
2. From the Clusters page, click the cluster that you want to add to the Favorites list.

3. On the Cluster details page, click the star icon.

Cluster favorite card

The Cluster favorite card enables you to view the capacity, configuration, and performance details of the individual clusters that you marked as favorites.

Cluster attributes

The Cluster favorite card displays the following attributes of individual clusters:

Cluster health status

An icon that indicates the health of the cluster. The possible values are Normal, Warning, Error, and Critical.

Cluster name

Name of the cluster.

Capacity

Total free space on the cluster.

Configuration

Configuration details of the cluster.

IP Address

IP address, or host name, of the cluster management logical interface (LIF) that was used to add the cluster.

Number of nodes

Number of nodes in the cluster.

Performance

Performance details of the cluster.

IOPS

Average number of I/O operations per second over the last 72 hours.

Throughput

Average throughput over the last 72 hours, in MBps .

Aggregate favorite card

The Aggregate favorite card enables you to view the capacity and performance details of the aggregates that you marked as favorites.

Aggregate attributes

The Aggregate favorite card displays the following aggregate attributes:

Aggregate health status

An icon that indicates the health of the aggregate. The possible values are Normal, Warning, Error, and Critical.

Aggregate name

Name of the aggregate.

Hierarchy

Hierarchy of the aggregate—the cluster to which the aggregate belongs.

Capacity

Percentage of free space available on the aggregate, and the estimated number of days until the aggregate becomes full.

Performance

Performance details of the aggregate.

IOPS

Average number of I/O operations per second over the last 72 hours.

Throughput

Average throughput over the last 72 hours, in MBps .

Latency

Average response time required for an operation, in milliseconds.

Volume favorite card

The Volume favorite card enables you to view the capacity, protection, and performance details of the volumes that you marked as favorites.

Volume attributes

The Volume favorite card displays the following volume attributes:

Volume health status

An icon that indicates the health status of the volume. The possible values are Normal, Warning, Error, and Critical.

Volume name

Name of the volume.

Hierarchy

Hierarchy of the volume; the cluster and SVM to which the volume belongs.

Capacity

Percentage of free space available on the volume, and the estimated number of days until the volume would become full.

Protection

Protection role that is set for the volume. The possible values are Unprotected, Not Applicable, Protected, and Destination.

Performance

Performance details for the volume.

IOPS

Average number of I/O operations per second over the last 72 hours.

Throughput

Average throughput over the last 72 hours, in MBps.

Latency

Average response time required for an operation, in milliseconds.

Monitoring performance

Unified Manager can be integrated with OnCommand Performance Manager, to give you access to performance data. OnCommand Performance Manager provides performance monitoring and incident root-cause analysis of systems running clustered Data ONTAP software. It is the performance management part of OnCommand Unified Manager.

Performance Manager helps you identify workloads that are over-using cluster components and decreasing the performance of other workloads on the cluster. It alerts you to these performance issues, called incidents, so that you can take corrective action and return performance back to normal. You can view and analyze incidents in the Performance Manager GUI or view them in the Unified Manager Dashboard.

Performance Manager is an analytics-based IT management software solution that helps you monitor the performance of FlexVol volumes, also called *user-defined workloads*, and internal workload activity, called *system-defined workloads*, on a system running clustered Data ONTAP. All monitored volumes must be in a QoS policy group. Infinite Volumes are not supported.

Performance Manager collects and analyzes workload activity to learn the expected range, or what it perceives to be normal activity for your environment. It uses the expected range and a dynamic performance threshold to monitor the I/O response time of each volume on a cluster. A high response time indicates which volumes are performing slower than normal. Slow response time also indicates when the performance of client applications that are using a volume has decreased. Performance Manager identifies the specific cluster component that is in contention, which is the location in the cluster where the performance issue lies. Performance Manager also provides a list of suggested actions you can take to try and address any performance issues yourself.

Performance Manager collects and analyzes workload activity every 5 minutes to detect performance incidents and detects configuration changes every 15 minutes. It retains a maximum of 90 days of historical performance and event data and forecasts the expected range for all monitored workloads.

Enabling performance monitoring

To enable performance monitoring from the Unified Manager web UI, you must perform two additional configuration tasks: creating a local user with limited Event Publisher role privileges and configuring connections between the Unified Manager server and one or more Performance Manager servers.

Note: When accessing the Performance Manager UI from a Unified Manager UI or email alert, you cannot access the Performance Manager Dashboard. Also, the Administration menu and Change Password option are not displayed.

Creating a user with Event Publisher role privileges

To support a connection between a Performance Manager server and Unified Manager, you must create a local user for Unified Manager and assign to it the Event Publisher role.

Before you begin

You must have the OnCommand Administrator role in Unified Manager.

About this task

When you configure a connection between a Performance Manager server and Unified Manager, the local user assigned the Event Publisher role is specified as the user under which performance event notification is posted in the Unified Manager web UI.

Steps

1. Log in to Unified Manager and navigate to the **Health** dashboard.
2. Click **Administration** > **Manage Users**.
3. In the **Manage Users** page, click **Add**.
4. In the **Add User** dialog box, select **Local User** for `type` and **Event Publisher** for `role`, and then enter the other required information.
5. Click **Add**.

After you finish

You can now configure a connection between one or more Performance Manager servers and Unified Manager.

Related concepts

Purpose of a connection between Performance Manager and Unified Manager on page 59

Related tasks

Configuring a partial integration connection between a Performance Manager server and Unified Manager on page 55

Configuring a full integration connection between a Performance Manager server and Unified Manager

To display performance issues discovered by a Performance Manager server in the Unified Manager web UI, you must configure a connection between Performance Manager and Unified Manager using the Performance Manager maintenance console.

Before you begin

- You intend to configure a full integration connection.
- The Unified Manager server must be installed with version 6.4 or later software.
- The versions of Unified Manager and Performance Manager must be compatible.
  The Interoperability Matrix contains the list of compatible versions.
  mysupport.netapp.com/matrix
- You must have a user ID authorized to log in to the maintenance console of the Performance Manager server.
- You must have the Performance Manager umadmin user ID.
- You must be prepared to specify the following information about the Unified Manager server:
  - Unified Manager server name or IP address
    If using an FQDN, the last part cannot be a single letter; for example, `vm.company.a` is invalid.
  - Unified Manager Administrator user name and password
  - Unified Manager Event Publisher user name and password

  Important: When Unified Manager is installed in a high-availability configuration, you must use the global IP address; you cannot use the Unified Manager server name or FQDN.

- The Unified Manager server, Performance Manager servers, and clusters that are being managed must be set to the same absolute (UTC) time (or they must use the same NTP server) or performance events are not correctly identified.

About this task

You can configure connections between a single Unified Manager server and up to five Performance Manager servers.

> **Important:** Implementing a full integration connection with a Unified Manager server cannot be undone. You cannot disable the connection to run the Performance Manager server in a stand-alone configuration.

Steps

1. Log in using SSH as the maintenance user to the Performance Manager host to access the maintenance console.
    - If Performance Manager is installed as a virtual appliance, log in as the maintenance user on the Performance Manager server.
    - If Performance Manager is installed on Red Hat Enterprise Linux, log in as umadmin (the maintenance user's automatically assigned name) to the Red Hat Enterprise Linux host machine.

    The Performance Manager maintenance console prompts are displayed.
2. Type the number of the menu option labeled **Unified Manager Integration**.
3. If prompted, enter the maintenance user password again.
4. Select **Full Integration > Enable Full Integration**.
5. When prompted, supply the Unified Manager server name or IP address (IPv4 or IPv6).

    The maintenance console checks the validity of the specified Unified Manager server name or IP address (IPv4 or IPv6) and Unified Manager server port and, if necessary, prompts you to accept the Unified Manager server trust certificate to support the connection.
6. When prompted, supply the Administrator user name and password.
7. When prompted, supply the Event Publisher user name and password.
8. When prompted, supply the unique name for this instance of Performance Manager.

    This name enables you to easily identify the Performance Manager server you want to manage when there are many instances integrated with Unified Manager.
9. Type **y** to confirm that the connection settings are correct, or type **n** if the settings are incorrect and you want to discard your changes.
10. If Performance Manager is installed on Red Hat Enterprise Linux, restart the Performance Manager server. If Performance Manager is installed as a virtual appliance, the virtual machine restarts automatically.

Result

After the connection is complete, all new performance events discovered by Performance Manager are reflected on the Unified Manager Dashboard page and Events page.

> **Note:** Until an initial performance event is discovered, the Unified Manager Dashboard page remains unchanged.

Configuring a partial integration connection between a Performance Manager server and Unified Manager

To display performance issues that are discovered by a Performance Manager server in the Unified Manager web UI, you must configure a partial integration connection between Performance Manager and Unified Manager in the Performance Manager maintenance console.

Before you begin

- You intend to configure a partial integration connection.
- The version of Unified Manager must be compatible with the version of Performance Manager. See the Interoperability Matrix for the list of compatible versions. *mysupport.netapp.com/matrix*
- You must have created a local user with Event Publisher privileges on Unified Manager server.
- You must have a user ID that is authorized to log in to the maintenance console of the Performance Manager server.
- You must have the Performance Manager umadmin user ID.
- You must have the following information:
    - Unified Manager server name or IP address
    - Unified Manager server port number (must be 443)
    - Event Publisher user name
    - Event Publisher password

    Note: When Unified Manager is installed in a high-availability configuration, you must use the global IP address; you cannot use the Unified Manager server name or fully qualified domain name (FQDN).
- The clusters that are to be managed by Performance Manager and Unified Manager must have been added to both Performance Manager and Unified Manager.
- The Unified Manager server, Performance Manager servers, and clusters that are being managed must be set to the same absolute (UTC) time (or they must use the same NTP server), or new performance events are not correctly identified.

About this task

You can configure connections between one Unified Manager server and up to five Performance Manager servers.

Steps

1. Log in using SSH as the maintenance user to the Performance Manager host.
    - If Performance Manager is installed as a virtual appliance, log in as the maintenance user to the maintenance console of the Performance Manager server.
    - If Performance Manager is installed on Red Hat Enterprise Linux, log in as umadmin (the maintenance user's automatically assigned name) to the Red Hat Enterprise Linux host machine.

    The Performance Manager maintenance console prompts are displayed.
2. Type the number of the menu option that is labeled **Unified Manager Integration**.

3. If prompted, enter the maintenance user password again.
4. Select **Partial Integration > Add Unified Manager Server Connection**.
5. When prompted, enter the Unified Manager server name or IP address (IPv4 or IPv6) and the Unified Manager server port information.

   The maintenance console checks the validity of the specified Unified Manager server name or IP address (IPv4 or IPv6), and the Unified Manager server port, and, if necessary, prompts you to accept the Unified Manager server trust certificate to support the connection. The default Unified Manager server port 443 must be used.
6. When prompted, enter the Event Publisher user name and password, and then confirm that the settings are correct.
7. If you want to configure an additional connection between the Unified Manager and another Performance Manager server, log in as the maintenance user to that Performance Manager server, and repeat Steps 2 through 4 for each connection.

Result

After the connection is complete, all new performance events that are discovered by Performance Manager are displayed on the Unified Manager Dashboard page and Events page.

> **Note:** Until an initial performance event is discovered, the Unified Manager Dashboard page remains unchanged.

Related concepts

Purpose of a connection between Performance Manager and Unified Manager on page 59

Changing Unified Manager connection settings

You can change the settings related to the Unified Manager server to which Performance Manager is connected, including the IP address, administrator name and password, and event publisher name and password. You can also refresh the connection details when a new Unified Manager security certificate has been generated and Performance Manager needs to accept the new certificate.

Before you begin

- You must have a user ID authorized to log in to the maintenance console of the Performance Manager server.
- You must have the Performance Manager umadmin user ID.
- Depending on the type of change you are making, you must be prepared to specify the following information about the Unified Manager server:
  - Unified Manager server name or IP address
  - Administrator user name and password
  - Event Publisher user name and password

Steps

1. Using SSH, log in as the maintenance user to the Performance Manager host.
   - If Performance Manager is installed as a virtual appliance, log in as the maintenance user to the maintenance console of the Performance Manager server.

- If Performance Manager is installed on Red Hat Enterprise Linux, log in as umadmin (the maintenance user's automatically assigned name) to the Red Hat Enterprise Linux machine that hosts the Performance Manager server.

   The Performance Manager maintenance console prompts are displayed.

2. Type the number of the menu option labeled **Unified Manager Integration**.
3. If prompted, enter the maintenance user password again.
4. Select **Full Integration** > **Modify Full Integration Settings**.
5. Change the settings as necessary.
6. Type **y** to confirm that the new connection settings are correct.
7. If Performance Manager is installed on Red Hat Enterprise Linux, restart the Performance Manager server. If Performance Manager is installed as a virtual appliance, the virtual machine restarts automatically.

Deleting a connection between a Performance Manager server and Unified Manager

If you no longer want to display performance issues that are discovered by a specific Performance Manager server in the Unified Manager web UI, you can delete the connection between that server and Unified Manager.

Before you begin

You must have credentials to log in to the maintenance console of the Performance Manager server.

You must have the umadmin credentials to log in to the Red Hat Enterprise Linux machine on which the Performance Manager server is installed.

About this task

The delete option is available only for a partial integration (Performance Event Publishing only) connection between Performance Manager and Unified Manager. You cannot delete a full integration connection between Performance Manager and Unified Manager.

If you are planning to delete a Performance Manager instance that has an existing connection to Unified Manager, you must delete the connection before deleting the instance.

Steps

1. Log in as the maintenance user to the maintenance console of the Performance Manager server.

   The Performance Manager maintenance console prompts are displayed.

2. Log in as the maintenance user (umadmin) to the Red Hat Enterprise Linux machine on which Performance Manager is installed.

   The Performance Manager maintenance console prompts are displayed.

3. In the maintenance console, type the number of the menu option that is labeled **Unified Manager Integration**.
4. If prompted, enter the maintenance user password again.
5. Select **Partial Integration** > **Delete Unified Manager Server Connection**.
6. When prompted whether you want to delete the connection, type **y** to delete the connection.

Result

Performance events that are discovered by the specific Performance Manager server are no longer displayed in the Unified Manager web UI.

Viewing details of performance incidents listed in the Dashboard

If the Unified Manager Dashboard page indicates performance issues and lists performance incidents, you can quickly view the details of those incidents to diagnose and possibly address the underlying causes of the incidents in question.

Before you begin

- Unified Manager must be connected with one or more Performance Manager servers.
- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must be able to log in to the Performance Manager server associated with the performance incident that you want to diagnose.

Steps

1. In the **Performance** pane of the **Dashboard** page, locate the performance incident that you want to view details of and click its hypertext link.

   This action invokes Performance Manager to display in a new tab on your browser.
2. If prompted, log in to Performance Manager.

   Performance Manager displays the details page for the performance incident that you want to examine.
3. Follow the procedures described in the Performance Manager help to further determine the cause of the performance incident.
4. When you are finished, click the OnCommand Unified Manager browser tab to return to the Unified Manager web UI.

Related concepts

Why a cluster component can be in contention on page 69
Types of workloads monitored by Performance Manager on page 60

Related tasks

Viewing details of performance incidents listed on the Events page on page 58

Related references

Performance event analysis and notification on page 66

Viewing details of performance incidents listed on the Events page

You can view the details of an unresolved performance incident listed on the Unified Manager Events page to diagnose the underlying cause of the incident in question.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- Unified Manager must be connected to one or more Performance Manager servers.

- You must be able to log in to the Performance Manager server associated with the performance incident that you want to diagnose.

Steps

1. Click **Events**.
2. In the **Filters** pane of **Events** page, locate the **Impact area** options and click **Performance**.

   All the unresolved performance incidents are listed together for you to survey.
3. In the Name column, locate and click the hypertext link for the performance incident in question.

   This action displays a new tab on your browser for the display of Performance Manager.
4. If prompted, log in to Performance Manager.

   Performance Manager displays the details page for the performance incident that you want to examine.
5. Follow the procedures described in the Performance Manager Help to further determine the cause of the performance incident.
6. When you have finished, click the OnCommand Unified Manager browser tab to return to the Unified Manager web UI.

Related concepts

Why a cluster component can be in contention on page 69
Types of workloads monitored by Performance Manager on page 60

Related tasks

Viewing details of performance incidents listed in the Dashboard on page 58

Related references

Performance event analysis and notification on page 66

Understanding more about performance monitoring

The performance monitoring and related configuration tasks that you complete in Unified Manager and Performance Manager are based on understanding what performance incidents are, what a connection between Unified Manager and Performance Manager enables, how performance incidents are displayed in Unified Manager, and how performance issues are discovered and analyzed in the connected Performance Manager application.

Purpose of a connection between Performance Manager and Unified Manager

A connection between a Performance Manager server and the Unified Manager server enables you to monitor through the Unified Manager web UI the performance issues that are detected by the Performance Manager server.

A connection between a Performance Manager server and the Unified Manager server is established through the menu option labeled "Unified Manager Server Connection" in the Performance Manager maintenance console.

Connections between multiple Performance Manager servers and Unified Manager

You can connect multiple Performance Manager servers (5 or fewer) to a single Unified Manager server. The multiple connections enable the Unified Manager operator to track the performance monitoring of multiple Performance Manager servers by viewing a single Unified Manager Dashboard page.

If you are connecting multiple Performance Manager servers to a single Unified Manager server, you must ensure that each monitored volume workload is not double-monitored, that is, monitored by more than one of the connected Performance Manager servers. This restriction prevents a single performance incident associated with a monitored volume workload from being displayed redundantly as multiple performance incidents on the Unified Manager Dashboard page.

Types of workloads monitored by Performance Manager

You can use Performance Manager to monitor the performance of two types of workloads: user-defined and system-defined.

User-defined workloads

The I/O throughput from applications to the cluster. These are processes involved in read and write requests. A FlexVol volume is a user-defined workload.

Note: Performance Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

Performance Manager requires that the volumes you want to monitor be in a QoS policy group:

- When using clustered Data ONTAP 8.3, a policy group is assigned to all volumes, either by the administrator or by Data ONTAP.
- When using clustered Data ONTAP 8.2.x, a policy group is assigned to all volumes, either by the administrator or by Performance Manager when the cluster is added to the UI. When Performance Manager analyzes the cluster for configuration changes every 15 minutes, it adds any new volumes not in a policy group to the default policy group.

Note: With clustered Data ONTAP 8.2.x, if an SVM, LUN, or File storage object is in a policy group, Performance Manager cannot monitor the volumes contained in that object and the overall analysis is impacted. You must remove the storage object from the policy group to correct this issue.

If one or more of the following is true for a workload, it cannot be monitored by Performance Manager:

- It is a data protection copy in read-only mode.
- It is an Infinite Volume.
- It is an offline data clone.
- It is a mirrored volume in a MetroCluster configuration.

System-defined workloads

The internal processes involved with storage efficiency, data replication, and system health, including:

- Storage efficiency, such as deduplication

- Disk health, which includes RAID reconstruct, disk scrubbing, and so on
- Data replication, such as SnapMirror copies
- Management activities
- File system health, which includes various WAFL activities
- File system scanners, such as WAFL scan
- Copy offload, such as offloaded storage efficiency operations from VMware hosts
- System health, such as volume moves, data compression, and so on
- Unmonitored volumes

Performance data for system-defined workloads is displayed in the GUI only when the cluster component used by these workloads is in contention. For example, you cannot search for the name of a system-defined workload to view its performance data in the GUI. If multiple system-defined workloads of the same type are displayed, a letter is appended to the workload name. The letter is intended for use by support personnel.

For more information about workloads in storage QoS, see the *System Administration Reference*.

Related concepts

Roles of workloads involved in a performance event on page 70
Understanding clusters and cluster objects on page 153
Why a cluster component can be in contention on page 69

Related references

Performance event analysis and notification on page 66

Workload performance measurement values

Performance Manager measures the performance of workloads on a cluster based on historical and expected statistical values, which form the expected range of values for the workloads. It compares the actual workload statistical values to the expected range to determine when workload performance is too high or too low. A workload that is not performing as expected triggers a performance event report to notify you.

In the following illustration, the actual value, in red, represents the actual performance statistics in the time frame. The actual value has crossed the performance threshold, which is the upper bounds of the expected range. The peak is the highest actual value in the time frame. The deviation measures the change between the expected values and the actual values, while the peak deviation indicates the largest change between the expected values and the actual values.

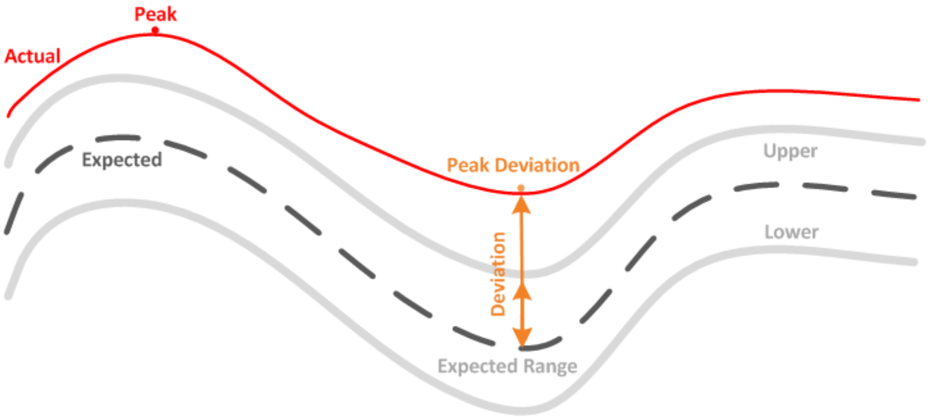

The following table lists the workload performance measurement values.

Measurement	Description
Activity	The percentage of the QoS limit used by the workloads in the policy group. **Note:** If Performance Manager detects a change to a policy group, such as adding or removing a volume or changing the QoS limit, the actual and expected values might exceed 100% of the set limit. If a value exceeds 100% of the set limit it is displayed as >100%. If a value is less than 1% of the set limit it is displayed as <1%.
Actual	The measured performance value at a specific time for a given workload.
Deviation	The change between the expected values and the actual values. It is the ratio of the actual value minus the expected value to the upper value of the expected range minus the expected value. **Note:** A negative deviation value indicates that workload performance is lower than expected, while a positive deviation value indicates that workload performance is higher than expected. If the expected values and the actual value are very low, in the hundredths or thousandths of a percent for example, the deviation will display N/A.
Expected	The expected values are based on the analysis of historical performance data for a given workload. Performance Manager analyzes these statistical values to determine the expected range of values.
Expected Range	The expected range of values is a forecast, or prediction, of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Performance Manager triggers a performance event alert.
Peak	The maximum value measured over a period of time.
Peak Deviation	The maximum deviation value measured over a period of time.
Queue Depth	The number of pending I/O requests that are waiting at the interconnect component.

Measurement	Description
Utilization	For the network processing, data processing, and aggregate components, the percentage of busy time to complete workload operations over a period of time. For example, the percentage of time for the network processing or data processing components to process an I/O request or for an aggregate to fulfill a read or write request.
Write Throughput	The amount of write throughput, in Megabytes per second (MBps), from workloads on a local cluster to the partner cluster in a MetroCluster configuration.

How Performance Manager uses workload latency to identify performance issues

The workload latency (response time) is the time it takes for a volume on a cluster to respond to I/O requests from client applications. Performance Manager uses the latency to detect and alert you to performance events.

A high latency means that requests from applications to a volume on a cluster are taking longer than usual. The cause of the high latency could be on the cluster itself, due to contention on one or more cluster components. High latency could also be caused by issues outside of the cluster, such as network bottlenecks, issues with the client hosting the applications, or issues with the applications themselves.

Note: Performance Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

Operations on the cluster, such as making backups or running deduplication, that increase their demand of cluster components shared by other workloads can also contribute to high latency. If the actual latency exceeds the performance threshold of the expected range, Performance Manager analyzes the event to determine whether it is a performance event that you might need to resolve. The latency is measured in milliseconds per operation (ms/op).

On the Volume Details page, you can view an analysis of the latency statistics to see how the activity of individual processes, such as read and write requests, compares to the overall latency statistics. The comparison helps you determine which operations have the highest activity or whether specific operations have abnormal activity that is impacting the latency for a volume. When analyzing performance events, you can use the latency statistics to determine whether an event was caused by an issue on the cluster. You can also identify the specific workload activities or cluster components that are involved in the event.

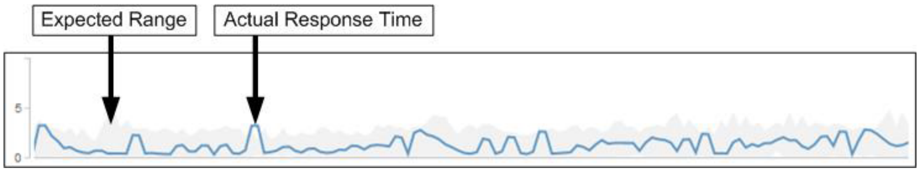

This example shows the Latency chart on the Volume Details page. The actual response time (latency) activity is a blue line and the expected range is gray.

Note: There can be gaps in the blue line if Performance Manager was unable to gather data. This can occur because the cluster or volume was unreachable, Performance Manager was turned off during that time, or the collection was taking longer than the 5 minute collection period.

Related concepts

What the expected range of performance is on page 64

How cluster operations can affect workload latency on page 64

Related references

Performance event analysis and notification on page 66

How cluster operations can affect workload latency

Operations (IOPS) represent the activity of all user-defined and system-defined workloads on a cluster. The IOPS statistics help you determine whether cluster processes, such as making backups or running deduplication, are impacting workload latency (response time) or might have caused, or contributed to, a performance event.

When analyzing performance events, you can use the IOPS statistics to determine whether a performance event was caused by an issue on the cluster. You can identify the specific workload activities that might have been the main contributors to the performance event. IOPS are measured in operations per second (ops/sec).

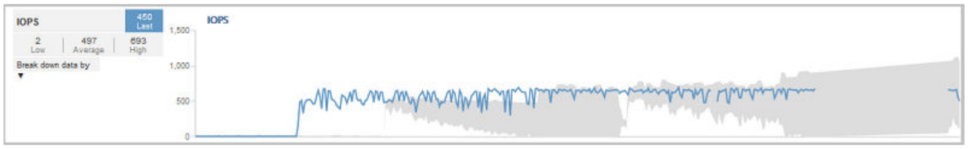

This example shows the IOPS chart on the Volume Details page. The actual operations statistics is a blue line and the expected range of operations statistics is gray.

Note: In some cases where a cluster is overloaded, Performance Manager might display the message `Data collection is taking too long on Cluster` *`cluster_name`*. This means that not enough statistics have been collected for Performance Manager to analyze. You need to reduce the resources the cluster is using so that statistics can be collected.

Related concepts

What the expected range of performance is on page 64

How Performance Manager uses workload latency to identify performance issues on page 63

What the expected range of performance is

The expected range of values is a forecast, or prediction, of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Performance Manager triggers a performance event alert.

For example, during regular business hours between 9:00 a.m. and 5:00 p.m., most employees might check their email between 9:00 a.m. and 10:30 a.m. The increased demand on the email servers means an increase in workload activity on the back-end storage during this time. Employees might notice slow response time from their email clients.

During the lunch hour between 12:00 p.m. and 1:00 p.m. and at the end of the work day after 5:00 p.m., most employees are likely away from their computers. The demand on the email servers typically decreases, also decreasing the demand on back-end storage. Alternatively, there could be scheduled workload operations, such as storage backups or virus scanning, that start after 5:00 p.m. and increase activity on the back-end storage.

Over several days, the increase and decrease in workload activity determines the expected range of activity, with upper and lower boundaries for a workload. When the actual workload activity for an object is outside the upper or lower boundaries, and remains outside the boundaries for a period of time, this might indicate that the object is being overused or underused.

How the expected range is formed

Performance Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the expected range for I/O response time and operations can be displayed in the GUI. The minimum required data collection does not account for all changes occurring from workload activity. After collecting the first 3 days of activity, Performance Manager adjusts the expected range, every 24 hours at 12:00 a.m., to reflect workload activity changes and establish a more accurate performance threshold.

> **Note:** Daylight savings time (DST) changes the system time, which alters the expected range of performance statistics for monitored workloads. Performance Manager immediately begins to correct the expected range, which takes approximately 15 days to complete. During this time you can continue to use Performance Manager, but, since Performance Manager uses the expected range to detect events, some events might not be accurate. Events detected prior to the time change are not affected. Manually changing the time on an ONTAP cluster, or on a Performance Manager server, to an earlier time will also affect the event analysis results.

Related concepts

What the expected range of performance is on page 64

Related references

How the expected range is used in performance analysis on page 65
Performance event analysis and notification on page 66
Workload performance measurement values on page 61

How the expected range is used in performance analysis

Performance Manager uses the expected range to represent the typical I/O latency (response time) and IOPS (operations) activity for your monitored workloads. It alerts you when the actual latency for a workload is above the upper bounds of the expected range, which triggers a performance event, so that you can analyze the performance issue and take corrective action for resolving it.

The expected range sets the performance baseline for the workload. Over time, Performance Manager learns from past performance measurements to forecast the expected performance and activity levels for the workload. The upper boundary of the expected range establishes the performance threshold. Performance Manager uses the baseline to determine when the actual latency or operations are above or below a threshold, or outside the bounds of their expected range. The comparison between the actual values and the expected values creates a performance profile for the workload.

When the actual latency for a workload exceeds the performance threshold, due to contention on a cluster component, the latency is high and the workload performs more slowly than expected. The performance of other workloads that share the same cluster components might also be slower than expected.

Performance Manager analyzes the threshold crossing event and determines whether the activity is a performance event. If the high workload activity remains consistent for a long period of time, such as several hours, Performance Manager considers the activity to be normal and dynamically adjusts the expected range to form the new performance threshold.

Some workloads might have consistently low activity, where the expected range for the operations or the latency does not have a high rate of change over time. To minimize the number of event alerts, during analysis of performance events, Performance Manager triggers an event only for low-activity volumes whose operations and latencies are much higher than expected.

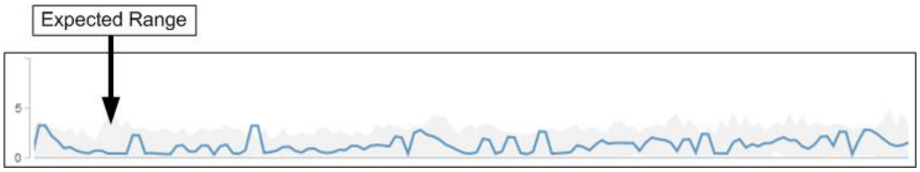

In this example, the latency for a volume has an expected range, in gray, of 0 milliseconds per operation (ms/op) at its lowest and 5 ms/op at its highest. If the actual latency, in blue, suddenly increases to 10 ms/op, due to an intermittent spike in network traffic or contention on a cluster component, it is then above the expected range and has exceeded the performance threshold.

When network traffic has decreased, or the cluster component is no longer in contention, the latency returns within the expected range. If the latency remains at or above 10 ms/op for a long period of time, you might need to take corrective action to resolve the event.

Related concepts

What the expected range of performance is on page 64

Related references

Performance event analysis and notification on page 66
Workload performance measurement values on page 61

Performance event analysis and notification

Performance events notify you about I/O performance issues on a volume workload caused by contention on a cluster component. Performance Manager analyzes the event to identify all workloads involved, the component in contention, and whether the event is still an issue that you might need to resolve.

Performance Manager monitors the I/O latency (response time) and IOPS (operations) for volumes on a cluster. When other workloads overuse a cluster component, for example, the component is in contention and cannot perform at an optimal level to meet workload demands. The performance of other workloads that are using the same component might be impacted, causing their latencies to increase. If the latency crosses the performance threshold, Performance Manager triggers a performance event and sends an email alert to notify you.

Event analysis

Performance Manager performs the following analyses, using the previous 15 days of performance statistics, to identify the victim workloads, bully workloads, and the cluster component involved in an event:

- Identifies victim workloads whose latency has crossed the performance threshold, which is the upper boundary of the expected range:
    - For volumes on HDD or Flash Pool (hybrid) aggregates, events are triggered only when the latency is greater than 5 milliseconds (ms) and the IOPS are more than 10 operations per second (ops/sec).
    - For volumes on all-SSD aggregates, events are triggered only when the latency is greater than 1 ms and the IOPS are more than 100 ops/sec.
- Identifies the cluster component in contention.

  Note: If the latency of victim workloads at the cluster interconnect is greater than 1 ms, Performance Manager treats this as significant and triggers an event for the cluster interconnect.

- Identifies the bully workloads that are overusing the cluster component and causing it to be in contention.
- Ranks the workloads involved, based on their deviation in utilization or activity of a cluster component, to determine which bullies have the highest change in usage of the cluster component and which victims are the most impacted.

An event might occur for only a brief moment and then correct itself after the component it is using is no longer in contention. A continuous event is one that reoccurs for the same cluster component within a five-minute interval and remains in the new state. For continuous events, Performance Manager triggers an alert after detecting the same event during two consecutive analysis intervals. Events that remain unresolved, which have a state of new, can display different description messages as workloads involved in the event change.

When an event is resolved, it remains available in Performance Manager as part of the record of past performance issues for a volume. Each event has a unique ID that identifies the event type and the volumes, cluster, and cluster components involved.

Note: A single volume can be involved in more than one event at the same time.

Event state

Events can be in one of the following states:

New

Indicates that the event is currently active. The issue causing the event has not corrected itself or has not been resolved. The performance counter for the cluster object remains above the performance threshold.

Obsolete

Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the cluster object is no longer above the performance threshold.

Event notification

The event alerts are displayed on the Dashboard, Volume Details page, and are sent to specified email addresses. If you have configured OnCommand Unified Manager to receive event alerts from Performance Manager, the events are also displayed on the Unified Manager Dashboard. You can view detailed analysis information about an event and get suggestions for resolving it on the Dynamic Threshold Event Details page.

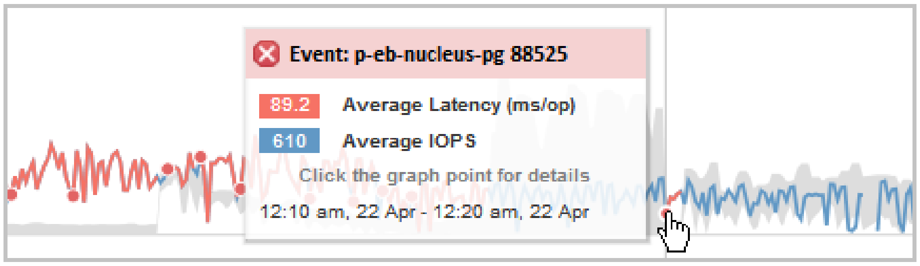

Event: p-eb-nucleus-pg 88525

89.2 Average Latency (ms/op)

610 Average IOPS

Click the graph point for details

12:10 am, 22 Apr - 12:20 am, 22 Apr

In this example, an event is indicated by a red dot (●) on the Latency chart on the Volume Details page. Hovering your mouse cursor over the red dot displays a popup with more details about the event and options for analyzing it.

Event interaction

On the Volume Details page, you can interact with events in the following ways:

- Moving the pointer over a red dot displays a message that shows the event ID, along with the latency, number of operations per second, and the date and time when the event was detected.

  If there are multiple events for the same time period, the message shows the number of events, along with the average latency and operations per second for the volume.
- Clicking a single event displays a dialog box that shows more detailed information about the event, including the cluster components that are involved, similar to the Summary section on the Dynamic Threshold Event Details page.

  The component in contention is circled and highlighted red. You can click either the event ID or **View full analysis** to view the full analysis on the Dynamic Threshold Event Details page. If there are multiple events for the same time period, the dialog box shows details about the three most recent events. You can click an event ID to view the event analysis on the Dynamic Threshold Event Details page. If there are more than three events for the same time period, clicking the red dot does not display the dialog box.

Related concepts

What the expected range of performance is on page 64
How Performance Manager determines the performance impact for an event on page 68
Roles of workloads involved in a performance event on page 70
Why a cluster component can be in contention on page 69
Types of workloads monitored by Performance Manager on page 60

Related tasks

Viewing details of performance incidents listed in the Dashboard on page 58
Viewing details of performance incidents listed on the Events page on page 58

How Performance Manager determines the performance impact for an event

Performance Manager uses the deviation in activity, utilization, write throughput, cluster component usage, or I/O latency (response time) for a workload to determine the level of impact to workload performance. This information determines the role of each workload in the event and how they are ranked on the Dynamic Threshold Event Details page.

Performance Manager compares the last analyzed values for a workload to the expected range of values. The difference between the values last analyzed and the expected range of values identifies the workloads whose performance was most impacted by the event.

For example, suppose a cluster contains two workloads: Workload A and Workload B. The expected range for Workload A is 5-10 milliseconds per operation (ms/op) and its actual latency is usually around 7 ms/op. The expected range for Workload B is 10-20 ms/op and its actual latency is usually around 15 ms/op. Both workloads are well within their expected range for latency. Due to contention on the cluster, the latency of both workloads increases to 40 ms/op, crossing the performance threshold, which is the upper bounds of the expected range, and triggering events. The deviation in latency, from the expected values to the values above the performance threshold, for Workload A is around 33 ms/op, and the deviation for Workload B is around 25 ms/op. The latency of both workloads spike to 40 ms/op, but Workload A had the bigger performance impact because it had the higher latency deviation at 33 ms/op.

On the Dynamic Threshold Event Details page, in the Workload Details table, you can sort workloads by their deviation in activity, utilization, or throughput for a cluster component. You can also sort workloads by latency. When you select a sort option, Performance Manager analyzes the deviation in activity, utilization, throughput, or latency since the event was detected from the expected values to determine the workload sort order. For the latency, the red dots (●) indicate a performance threshold crossing by a victim workload, and the subsequent impact to the latency. Each red dot indicates a higher level of deviation in latency, which helps you identify the victim workloads whose latency was impacted the most by an event.

Related concepts

What the expected range of performance is on page 64
Roles of workloads involved in a performance event on page 70
Why a cluster component can be in contention on page 69
Types of workloads monitored by Performance Manager on page 60

Related tasks

Viewing details of performance incidents listed in the Dashboard on page 58
Viewing details of performance incidents listed on the Events page on page 58

Related references

Performance event analysis and notification on page 66

Why a cluster component can be in contention

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

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

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

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

Network

Represents the wait time of I/O requests by the iSCSI protocols or the Fibre Channel (FC) protocols on the cluster. The wait time is time spent waiting for iSCSI Ready to Transfer (R2T) or FCP Transfer Ready (XFER_RDY) transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the block protocol layer is impacting the latency of one or more workloads.

Network Processing

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

Policy Group

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

Cluster Interconnect

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

Data Processing

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

MetroCluster Resources

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

Aggregate or SSD Aggregate

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

Related concepts

What the expected range of performance is on page 64
Roles of workloads involved in a performance event on page 70
Types of workloads monitored by Performance Manager on page 60

Related tasks

Viewing details of performance incidents listed in the Dashboard on page 58
Viewing details of performance incidents listed on the Events page on page 58

Related references

Performance event analysis and notification on page 66

Roles of workloads involved in a performance event

Performance Manager uses roles to identify the involvement of a workload in a performance event. The roles include victims, bullies, and sharks. A user-defined workload can be a victim, bully, and shark at the same time.

The following table defines the workload roles:

Role	Description
Victim	A user-defined workload whose performance has decreased due to other workloads, called bullies, that are over-using a cluster component. Only user-defined workloads are identified as victims. Performance Manager identifies victim workloads based on their deviation in latency, where the actual latency, during an event, has greatly increased from its expected range of latency.

Role	Description
Bully	A user-defined or system-defined workload whose over-use of a cluster component has caused the performance of other workloads, called victims, to decrease. Performance Manager identifies bully workloads based on their deviation in usage of a cluster component, where the actual usage, during an event, has greatly increased from its expected range of usage.
Shark	A user-defined workload with the highest usage of a cluster component compared to all workloads involved in an event. Performance Manager identifies shark workloads based on their usage of a cluster component during an event.

Workloads on a cluster can share many of the cluster components, such as storage aggregates and the CPU for network and data processing. When a workload, such as a volume, increases its usage of a cluster component to the point that the component cannot efficiently meet workload demands, the component is in contention. The workload that is over-using a cluster component is a bully. The other workloads that share those components, and whose performance is impacted by the bully, are the victims. Activity from system-defined workloads, such as deduplication or Snapshot copies, can also escalate into “bullying”.

When Performance Manager detects an event, it identifies all workloads and cluster components involved, including the bully workloads that caused the event, the cluster component that is in contention, and the victim workloads whose performance has decreased due to the increased activity of bully workloads.

Note: If Performance Manager cannot identify the bully workloads, it only alerts on the victim workloads and the cluster component involved.

Performance Manager can identify workloads that are victims of bully workloads, and also identify when those same workloads become bully workloads. A workload can be a bully to itself. For example, a high-performing workload that is being throttled by a policy group limit causes all workloads in the policy group to be throttled, including itself. A workload that is a bully or a victim in an ongoing performance event might change its role or no longer be a participant in the event. On the Volume Details page, in the Events List table, when the selected volume changes its participant role, the date and time of the role change is displayed.

Related concepts

Understanding clusters and cluster objects on page 153
Why a cluster component can be in contention on page 69
Types of workloads monitored by Performance Manager on page 60

Related tasks

Viewing details of performance incidents listed in the Dashboard on page 58
Viewing details of performance incidents listed on the Events page on page 58

Related references

Performance event analysis and notification on page 66

Configuration changes detected by Performance Manager

Performance Manager monitors your clusters for configuration changes to help you determine whether a change might have caused or contributed to a performance event. The Performance

Explorer page and the Volume Details page display a change event icon (●) to indicate the date and time when the change was detected.

You can review the performance charts in the Performance Explorer page and in the Volume Details page to see whether the change event impacted the performance of the selected cluster object. If the change was detected at or around the same time as a performance event, the change might have contributed to the issue, which caused the event alert to trigger.

Performance Manager can detect the following change events, which are categorized as Informational events:

- A volume moves between aggregates.
  Performance Manager can detect when the move is in progress, completed, or failed. If Performance Manager is down during a volume move, when Performance Manager is back up it detects the volume move and displays a change event for it.
- The throughput (MBps) limit of a QoS policy group that contains one or more monitored workloads changes.
  Changing a policy group limit can cause intermittent spikes in the latency (response time), which might also trigger events for the policy group. The latency gradually returns back to normal and any events caused by the spikes become obsolete.
- A node in an HA pair takes over or gives back the storage of its partner node.
  Performance Manager can detect when the takeover, partial takeover, or giveback operation has been completed. If the takeover is caused by a panicked node, Performance Manager does not detect the event.
- A Data ONTAP upgrade or revert operation is completed successfully.
  The previous version and new version are displayed.

Related concepts

How the maximum throughput limit works on page 164
What an HA pair is on page 160

Managing events and alerts

Events help you to identify issues in the clusters that are monitored. You can configure alerts to send notification automatically when specific events or events of certain severity types occur.

What events are

Events are notifications that are generated automatically when a predefined condition occurs or when an object crosses a threshold. These events enable you to take action to prevent issues that can lead to poor performance and system unavailability. Events include an impact area, severity, and impact level.

Events are categorized by the type of impact area such as availability, capacity, configuration, or protection. Events are also assigned a severity type and impact level that assist you in determining if immediate action is required.

You can configure alerts to send notification automatically when specific events or events of a specific severity occur.

Obsolete, resolved, and informational events are automatically logged and retained for a default of 180 days.

It is important that you take immediate corrective action for events with severity level Error or Critical.

Related concepts

What alerts are on page 102

Related references

Description of event severity types on page 81
Description of event impact levels on page 82
Description of event impact areas on page 82

What performance events are

Performance events are incidents or configuration changes related to workload performance on a cluster. They help you identify workloads with slow response times and cluster configuration changes that might have caused or contributed to the slow response times.

When Performance Manager detects multiple occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events.

Types of performance events

Performance Manager can detect the following types of performance events:

User-defined threshold policy events

Performance issues based on custom threshold values that you have set. You configure threshold policies for cluster objects; for example, aggregates and volumes, so that events are generated when a threshold value for a performance counter has been breached.

You must define a threshold policy and assign it to a cluster object to receive these events.

System-defined threshold policy events

Performance issues based on threshold values that are system-defined. These threshold policies are provided with the installation of Performance Manager to cover common performance problems.

These threshold policies are enabled by default, and you might see events shortly after adding a cluster.

Dynamic threshold events

Performance issues that are the result of failures or errors in an IT infrastructure, or from workloads overutilizing cluster resources. The cause of these events might be a simple issue that corrects itself over a period of time or can be addressed with a repair or configuration change. A dynamic threshold event indicates that volume workloads on a system running ONTAP are slow due to other workloads with high usage of shared cluster components.

These thresholds are enabled by default, and you might see events after three days of collecting data from a new cluster.

Configuration change events

A change event is generated when a logical or physical storage object in a system running ONTAP is added, modified, or deleted. A change that impacts the performance of one or more workloads might be a contributor to an event.

Related concepts

What the expected range of performance is on page 64
Configuration changes detected by Performance Manager on page 71
Roles of workloads involved in a performance event on page 70
Types of workloads monitored by Performance Manager on page 60

Related references

Performance event analysis and notification on page 66

What Event Management System events are

The Event Management System (EMS) collects event data from different parts of the Data ONTAP kernel, and provides event forwarding mechanisms. These ONTAP events are captured and reported as EMS events in Unified Manager.

The EMS events that are generated by Unified Manager are different from other events in Unified Manager. EMS events have the following characteristics:

- The name of an EMS event is in a dot-notation format, and includes a collection of named attributes.
- The life cycle of an EMS event is one monitoring cycle.
  The time to obsolete an EMS event for a cluster depends on the refresh of that particular cluster. The EMS event is made obsolete only after the cluster is refreshed, after 15 minutes, from when the event was created.
- An EMS event is reported as soon as the event occurs in Data ONTAP.

The subscribed EMS events are validated, and only validated events are applied to clusters in Unified Manager. You can configure Unified Manager alerts for the reported EMS events.

Configuring event settings

You can specify the number of days an event is retained in the Unified Manager server before it is automatically deleted. Only events that are resolved, obsolete, or of type Information are deleted. You can also specify the frequency with which these events are deleted or you can also manually delete the events.

Before you begin

You must have the OnCommand Administrator role to change the event settings.

About this task

Retaining events for more than 180 days affects the server performance and is not recommended. The lower limit for the event retention period is 7 days; there is no upper limit.

Steps

1. Click **> Health** > **Administration** > **Manage Events**.
2. In the **Manage Events** page, click **Event Retention Settings**.
3. Configure the appropriate settings in the **Event Retention Settings** dialog box.
4. Click **Save and Close**.

Related tasks

Adding a user on page 379

Configuring notification settings

You can configure the settings for the Unified Manager server to send alert notifications when an event is generated or when it is assigned to a user. You can configure the corresponding mail server to be used and various notification mechanisms—for example, alert notifications can be sent as emails or SNMP traps.

Before you begin

The following information must be available:

- Email address from which the alert notification is sent
- Host name, user name, password, and default port to configure the SMTP server
- SNMP version, trap destination host, outbound trap port, and community to configure the SNMP trap

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **General Settings** > **Notification**.

4. **4.** Configure the appropriate settings.

   You can specify the email address and SMTP server from which the alert notifications are sent, and enter the SNMP trap settings.

   Tip: If the host name of the SMTP server cannot be resolved, you can specify the IP address (IPv4 or IPv6) of the SMTP server instead of the host name.

Related tasks

Adding a user on page 379
Adding an alert on page 102

Subscribing to ONTAP EMS events

You can subscribe to Event Management System (EMS) events that systems installed with ONTAP software report. EMS events are captured by OnCommand Unified Manager only if you have subscribed to these events. These captured events are validated and applied to clusters in Unified Manager.

About this task

You can configure alerts for the ONTAP EMS events to which you subscribe, and you can create custom scripts to be executed for these events.

You can subscribe to any number of EMS events. All the events to which you subscribe are validated, and only the validated events are applied to the cluster in Unified Manager.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Events**.
3. Click **Subscribe to EMS events**.
4. In the **Subscribe to EMS events** dialog box, enter the name of the ONTAP EMS event to which you want to subscribe.

   To view the names of the EMS message events to which you can subscribe, from the ONTAP cluster shell, use the `event route show` command (prior to ONTAP 9) or the `event catalog show` command (ONTAP 9 or later).
5. Click **Add**.

   The EMS event is added to the Subscribed EMS events list, but the status “Unknown” appears in the Applicable to Cluster column.
6. Click **Save and Close** to register the EMS event subscription with the cluster.
7. Click **Subscribe to EMS events**.

   The status “Yes” appears in the Applicable to Cluster column for the EMS event you added.

   If the status is not “Yes”, check the spelling of the ONTAP EMS event. If the name is entered incorrectly, you must remove the incorrect event and then add the event again.

Result

Note: If you do not receive the ONTAP EMS events to which you have subscribed, there might be an issue with the DNS configuration of the cluster which is blocking the cluster from reaching the

Unified Manager server. In this occurs, the cluster administrator must correct the DNS configuration on the cluster and then restart Unified Manager. This will flush the pending EMS events to the Unified Manager server.

Viewing event details

You can view details about an event that is triggered by Unified Manager to take corrective action. For example, if there is an event Volume Offline, you can click that event to view the details and perform corrective actions.

Before you begin

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

About this task

The event details include information such as the source of the event, cause of the event, and any notes related to the event.

Steps

1. Click **Events**.
2. In the **Events** page, click the event name for which you want to view the details.

   The event details are displayed in the Event details page.

Related tasks

Adding a user on page 379
Assigning events on page 78
Adding an alert on page 102
Acknowledging and resolving events on page 80
Adding and reviewing notes about an event on page 79

Disabling or enabling events

You can disable events globally to prevent the generation of notifications for events that are not important in your environment. You can enable events that are disabled when you want to resume receiving notifications for them.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

When you disable an event for an object (for example, the vol offline event), and you enable the disabled event later, Unified Manager does not generate new events for objects that went offline when the event was in the disabled state. Unified Manager generates a new event only when there is a change in the object state after the event was reenabled.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Events**.
3. Disable or enable events by choosing one of the following options:

If you want to...	Then do this...
Disable events	**a.** Click **Disable**. **b.** In the Disable Events dialog box, select the events that you want to disable based on the event severity. **c.** Click **Save and Close**. **d.** Verify that the events that you disabled are displayed in the list view of the Manage Events page.
Enable events	**a.** In the list view of the Manage Events page, select the check box for the event, or events, that you want to enable. **b.** Click **Enable**.

Assigning events

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

Before you begin

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

Steps

1. Click **Events**.
2. From the events list on the **Events** page, select one or more events that you want to assign.
3. Assign the event by choosing one of the following options:

If you want to assign the event to...	Then do this...
Yourself	Click **Assign To** > **Me**.
Another user	**a.** Click **Assign To** > **Another user**. **b.** In the Assign Owner dialog box, enter the user name, or select a user from the drop-down list. **c.** Click **Assign**. An email notification is sent to the user. **Note:** If you do not enter a user name or select a user from the drop-down list, and click **Assign**, the event remains unassigned.

Related tasks

Adding a user on page 379

Viewing unassigned events

You can view unassigned events and then assign each of them to a user who can resolve them.

Before you begin

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

Steps

1. Click **Events**.

   By default, New and Acknowledged events are displayed on the Events page.
2. From the **Filters** pane, select the **Unassigned** filter option in the **Assigned To** area.

Related tasks

Adding a user on page 379
Assigning events on page 78

Adding and reviewing notes about an event

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

Before you begin

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

Steps

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

Related tasks

Adding a user on page 379

Acknowledging and resolving events

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

Before you begin

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

About this task

You can acknowledge and resolve multiple events simultaneously.

Steps

1. Click **Events**.
2. From the events list, perform the following actions to acknowledge the events:

If you want to...	Do this...
Acknowledge and mark a single event as resolved	**a.** Click the event name. **b.** From the Event details page, determine the cause of the event. **c.** Click **Acknowledge**. **d.** Take appropriate corrective action. **e.** Click **Mark As Resolved**.
Acknowledge and mark multiple events as resolved	**a.** Determine the cause of the events from the respective Event details page. **b.** Select the events. **c.** Click **Acknowledge**. **d.** Take appropriate corrective actions. **e.** Click **Mark As Resolved**.

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

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

Related tasks

Adding a user on page 379
Adding and reviewing notes about an event on page 79

Understanding more about events

Understanding the concepts about events helps you to manage your clusters and cluster objects efficiently and to define alerts appropriately.

Event state definitions

The state of an event helps you identify whether an appropriate corrective action is required. An event can be New, Acknowledged, Resolved, or Obsolete.

The different event states are as follows:

New

The state of a new event.

Acknowledged

The state of an event when you have acknowledged it.

Resolved

The state of an event when it is marked as resolved.

Obsolete

The state of an event when it is automatically corrected or when the cause of the event is no longer valid.

Note: You cannot acknowledge or resolve an obsolete event.

Example of different states of an event

The following examples illustrate the manual and automatic event state changes.

When the event Cluster Not Reachable is triggered, the event state is New. When you acknowledge the event, the event state changes to Acknowledged. When you have taken an appropriate corrective action, you must mark the event as resolved. The event state then changes to Resolved.

If the Cluster Not Reachable event is generated due to a power outage, after the power is back the cluster starts functioning without any administrator intervention. Therefore, the Cluster Not Reachable event is no longer valid, and the event state changes to Obsolete in the next monitoring cycle.

Unified Manager sends an alert when an event is in the Obsolete or Resolved state. The email subject line and email content of an alert provides information about the event state. An SNMP trap too includes information about the event state.

Description of event severity types

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

Critical

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

Error

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

Warning

The event source experienced an occurrence that you should be aware of. Events of this severity do not cause service disruption, and immediate corrective action might not be required.

Information

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

Related references

List of events and severity types on page 83

Description of event impact levels

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

Incident

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

Risk

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

Event

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

Description of event impact areas

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

Availability

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

Capacity

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

Configuration

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

Performance

Performance events, also called incidents, notify you of resource, configuration, or activity conditions on your storage cluster that might adversely affect the speed of data storage input or retrieval on your monitored SVM and volumes. Description of performance impact events is provided in OnCommand Performance Manager help.

Protection

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

How object status is computed

Object status is determined by the most severe event that currently holds a New or Acknowledged state. For example, if an object status is Error, then one of the object's events has a severity type of Error. When corrective action has been taken, the event state moves to Resolved.

Related concepts

Event state definitions on page 81

Related tasks

Acknowledging and resolving events on page 80

List of events and severity types

You can use the list of events to become more familiar with event categories, event names, and the severity type of each event that you might see in OnCommand Unified Manager. Events are listed in alphabetical order by object category.

Aggregate events

Aggregate events provide you with information about the status of aggregates so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Aggregate Offline (ocumEvtAggregateStateOffline)	Incident	Aggregate	Critical
Aggregate Failed (ocumEvtAggregateStateFailed)	Incident	Aggregate	Critical
Aggregate Restricted (ocumEvtAggregateStateRestricted)	Risk	Aggregate	Warning
Aggregate Reconstructing (ocumEvtAggregateRaidStateReconstructing)	Risk	Aggregate	Warning
Aggregate Degraded (ocumEvtAggregateRaidStateDegraded)	Risk	Aggregate	Warning
MetroCluster Aggregate Left Behind (ocumEvtMetroClusterAggregateLeftBehind)	Risk	Aggregate	Error
MetroCluster Aggregate Mirroring Degraded (ocumEvtMetroClusterAggregateMirrorDegraded)	Risk	Aggregate	Error

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Aggregate Space Nearly Full (ocumEvtAggregateNearlyFull)	Risk	Aggregate	Warning
Aggregate Space Full (ocumEvtAggregateFull)	Risk	Aggregate	Error
Aggregate Days Until Full (ocumEvtAggregateDaysUntilFullSoon)	Risk	Aggregate	Error
Aggregate Overcommitted (ocumEvtAggregateOvercommitted)	Risk	Aggregate	Error
Aggregate Nearly Overcommitted (ocumEvtAggregateAlmostOvercommitted)	Risk	Aggregate	Warning
Aggregate Snapshot Reserve Full (ocumEvtAggregateSnapReserveFull)	Risk	Aggregate	Warning
Aggregate Growth Rate Abnormal (ocumEvtAggregateGrowthRateAbnormal)	Risk	Aggregate	Warning

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Aggregate Discovered (Not applicable)	Event	Aggregate	Information
Aggregate Renamed (Not applicable)	Event	Aggregate	Information
Aggregate Deleted (Not applicable)	Event	Node	Information

Cluster events

Cluster events provide information about the status of clusters which enables you to monitor the clusters for potential problems. The events are grouped by impact area, and include the event name, trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Cluster Lacks Spare Disks (ocumEvtDisksNoSpares)	Risk	Cluster	Warning
Cluster Not Reachable (ocumEvtClusterUnreachable)	Risk	Cluster	Error

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
Cluster Monitoring Failed (ocumEvtClusterMonitoringFailed)	Risk	Cluster	Warning
MetroCluster Spare Disks Left Behind (ocumEvtSpareDiskLeftBehind)	Risk	Cluster	Error
MetroCluster Automatic Unplanned Switchover Disabled (ocumEvtMccAutomaticUnplannedSwitchOverDisabled)	Risk	Cluster	Warning

Impact area: configuration

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
Node Added (Not applicable)	Event	Cluster	Information
Node Removed (Not applicable)	Event	Cluster	Information
Cluster Removed (Not applicable)	Event	Cluster	Information
Cluster Add Failed (Not applicable)	Event	Cluster	Error
Cluster Name Changed (Not applicable)	Event	Cluster	Information
Critical EMS received (Not applicable)	Event	Cluster	Critical
Alert EMS received (Not applicable)	Event	Cluster	Error
Error EMS received (Not applicable)	Event	Cluster	Warning
Warning EMS received (Not applicable)	Event	Cluster	Warning
Debug EMS received (Not applicable)	Event	Cluster	Warning
Notice EMS received (Not applicable)	Event	Cluster	Warning
Informational EMS received (Not applicable)	Event	Cluster	Warning

Note: ONTAP EMS events are categorized into three Unified Manager event severity levels.

Unified Manager event severity level	ONTAP EMS event severity level
Critical	Emergency Critical
Error	Alert
Warning	Error Warning Debug Notice Informational

Port events

Port events provide you with status about cluster ports so that you can monitor changes or problems on the port, like whether the port is down.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Port Status Down (ocumEvtPortStatusDown)	Incident	Node	Critical

Disks events

Disks events provide you with information about the status of disks so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Flash Disks - Spare Blocks Almost Consumed (ocumEvtClusterFlashDiskFewerSpareBlockError)	Risk	Cluster	Error
Flash Disks - No Spare Blocks (ocumEvtClusterFlashDiskNoSpareBlockCritical)	Incident	Cluster	Critical
Some Unassigned Disks (ocumEvtClusterUnassignedDisksSome)	Risk	Cluster	Warning
Some Failed Disks (ocumEvtDisksSomeFailed)	Incident	Cluster	Critical

Enclosures events

Enclosures events provide you with information about the status of disk shelf enclosures in your data center so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
Disk Shelf Fans Failed (ocumEvtShelfFanFailed)	Incident	Storage shelf	Critical
Disk Shelf Power Supplies Failed (ocumEvtShelfPowerSupplyFailed)	Incident	Storage shelf	Critical
Disk Shelf Multipath Not Configured (ocumDiskShelfConnectivityNotInMultiPath) This event does not apply to: • Clusters that are in a MetroCluster configuration • The following platforms: FAS2554, FAS2552, FAS2520, and FAS2240	Risk	Node	Warning
Disk Shelf Path Failure (ocumDiskShelfConnectivityPathFailure)	Risk	Storage Shelf	Warning

Impact area: configuration

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
Disk Shelf Discovered (Not applicable)	Event	Node	Information
Disk Shelves Removed (Not applicable)	Event	Node	Information

Fans events

Fans events provide you with information about the status fans on nodes in your data center so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
One or More Failed Fans (ocumEvtFansOneOrMoreFailed)	Incident	Node	Critical

Flash card events

Flash card events provide you with information about the status of the flash cards installed on nodes in your data center so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
Flash Cards Offline (ocumEvtFlashCardOffline)	Incident	Node	Critical

Inodes events

Inode events provide information when the inode is full or nearly full so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
Inodes Nearly Full (ocumEvtInodesAlmostFull)	Risk	Volume	Warning
Inodes Full (ocumEvtInodesFull)	Risk	Volume	Error

Logical interface events

Logical interface events provide information about the status of your LIFs, so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
LIF Status Down (ocumEvtLifStatusDown)	Risk	Interface	Error
LIF Failover Not Possible (ocumEvtLifFailoverNotPossible)	Risk	Interface	Warning
LIF Not At Home Port (ocumEvtLifNotAtHomePort)	Risk	Interface	Warning

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
LIF Route Not Configured (Not applicable)	Event	Interface	Information

LUN events

LUN events provide you with information about the status of your LUNs, so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
LUN Offline (ocumEvtLunOffline)	Incident	LUN	Critical
Single Active Path To Access LUN (ocumEvtLunSingleActivePath)	Risk	LUN	Warning
No Active Paths To Access LUN (ocumEvtLunNotReachable)	Incident	LUN	Critical
No Optimized Paths To Access LUN (ocumEvtLunOptimizedPathInactive)	Risk	LUN	Warning
No Paths To Access LUN From HA Partner (ocumEvtLunHaPathInactive)	Risk	LUN	Warning

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Insufficient Space For LUN Snapshot Copy (ocumEvtLunSnapshotNotPossible)	Risk	Volume	Warning

MetroCluster Bridge events

MetroCluster Bridge events provide you with information about the status of the bridges so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Bridge Unreachable (ocumEvtBridgeUnreachable)	Incident	MetroCluster Bridge	Critical

Event name (Trap name)	Impact level	Source type	Severity
Bridge Temperature Abnormal (ocumEvtBridgeTemperatureAbnormal)	Incident	MetroCluster Bridge	Critical

MetroCluster Connectivity events

Connectivity events provide you with information about the connectivity between the components of a cluster and between clusters in a MetroCluster configuration so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
All Inter-Switch Links Down (ocumEvtMetroClusterAllISLBetweenSwitchesDown)	Incident	MetroCluster inter-switch connection	Critical
All Links Between MetroCluster Partners Down (ocumEvtMetroClusterAllLinksBetweenPartnersDown)	Incident	MetroCluster relationship	Critical
FC-SAS Bridge To Storage Stack Link Down (ocumEvtBridgeSasPortDown)	Incident	MetroCluster bridge stack connection	Critical
MetroCluster Configuration Switched Over ((ocumEvtMetroClusterDRStatusImpacted)	Risk	MetroCluster relationship	Warning
MetroCluster Configuration Partially Switched Over (ocumEvtMetroClusterDRStatusPartiallyImpacted)	Risk	MetroCluster relationship	Error
MetroCluster Disaster Recovery Capability Impacted (ocumEvtMetroClusterDRStatusImpacted)	Risk	MetroCluster relationship	Critical
MetroCluster Partners Not Reachable Over Peering Network (ocumEvtMetroClusterPartnersNotReachableOverPeeringNetwork)	Incident	MetroCluster relationship	Critical
Node To FC Switch All FC-VI Interconnect Links Down (ocumEvtMccNodeSwitchFcviLinksDown)	Incident	MetroCluster node switch connection	Critical
Node To FC Switch One Or More FC-Initiator Links Down (ocumEvtMccNodeSwitchFcLinksOneOrMoreDown)	Risk	MetroCluster node switch connection	Warning
Node To FC Switch All FC-Initiator Links Down (ocumEvtMccNodeSwitchFcLinksDown)	Incident	MetroCluster node switch connection	Critical
Switch To FC-SAS Bridge FC Link Down (ocumEvtMccSwitchBridgeFcLinksDown)	Incident	MetroCluster switch bridge connection	Critical

Event name (Trap name)	Impact level	Source type	Severity
Inter Node All FC VI InterConnect Links Down (ocumEvtMccInterNodeLinksDown)	Incident	Inter-node connection	Critical
Inter Node One Or More FC VI InterConnect Links Down (ocumEvtMccInterNodeLinksOneOrMoreDown)	Risk	Inter-node connection	Warning
Node To Bridge Link Down (ocumEvtMccNodeBridgeLinksDown)	Incident	Node bridge connection	Critical
Node to Storage Stack All SAS Links Down (ocumEvtMccNodeStackLinksDown)	Incident	Node stack connection	Critical
Node to Storage Stack One Or More SAS Links Down (ocumEvtMccNodeStackLinksOneOrMoreDown)	Risk	Node stack connection	Warning

MetroCluster switch events

MetroCluster switch events provide you with information about the status of the MetroCluster switches so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Switch Temperature Abnormal (ocumEvtSwitchTemperatureAbnormal)	Incident	MetroCluster Switch	Critical
Switch Unreachable (ocumEvtSwitchUnreachable)	Incident	MetroCluster Switch	Critical
Switch Fans Failed (ocumEvtSwitchFansOneOrMoreFailed)	Incident	MetroCluster Switch	Critical
Switch Power Supplies Failed (ocumEvtSwitchPowerSuppliesOneOrMoreFailed)	Incident	MetroCluster Switch	Critical
Switch Temperature Sensors Failed (ocumEvtSwitchTemperatureSensorFailed) **Note:** This event is applicable only for Cisco switches.	Incident	MetroCluster Switch	Critical

Node events

Node events provide you with information about node status so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Node Root Volume Space Nearly Full (ocumEvtClusterNodeRootVolumeSpaceNearlyFull)	Risk	Node	Warning

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Node Renamed (Not applicable)	Event	Node	Information

NVRAM battery events

NVRAM battery events provide you with information about the status of your batteries so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
NVRAM Battery Low (ocumEvtNvramBatteryLow)	Risk	Node	Warning
NVRAM Battery Discharged (ocumEvtNvramBatteryDischarged)	Risk	Node	Error
NVRAM Battery Overly Charged (ocumEvtNvramBatteryOverCharged)	Incident	Node	Critical

Power supplies events

Power supplies events provide you with information about the status of your hardware so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
One or More Failed Power Supplies (ocumEvtPowerSupplyOneOrMoreFailed)	Incident	Node	Critical

Protection events

Protection events tell you if a job has failed or been aborted so that you can monitor for problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: protection

Event name (Trap name)	Impact level	Source type	Severity
Protection Job Failed (ocumEvtProtectionJobTaskFailed)	Incident	Volume or storage service	Critical
Protection Job Aborted (ocumEvtProtectionJobAborted)	Risk	Volume or storage service	Warning

Qtree events

Qtree events provide you with information about the qtree capacity and the file and disk limits so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Qtree Files Hard Limit Reached (ocumEvtQtreeFilesHardLimitReached)	Incident	Qtree	Critical
Qtree Files Soft Limit Breached (ocumEvtQtreeFilesSoftLimitBreached)	Risk	Qtree	Warning
Qtree Space Hard Limit Reached (ocumEvtQtreeSpaceHardLimitReached)	Incident	Qtree	Critical
Qtree Space Soft Limit Breached (ocumEvtQtreeSpaceSoftLimitBreached)	Risk	Qtree	Warning

Service processor events

Service processor events provide you with information about the status of your processor so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Service Processor Not Configured (ocumEvtServiceProcessorNotConfigured)	Risk	Node	Warning
Service Processor Offline (ocumEvtServiceProcessorOffline)	Risk	Node	Error

SnapMirror relationship events

SnapMirror relationship events provide you with information about the status of your SnapMirror relationships so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: protection

Event name (Trap name)	Impact level	Source type	Severity
SnapMirror Relationship Unhealthy (ocumEvtSnapmirrorRelationshipUnhealthy)	Risk	SnapMirror relationship	Warning
SnapMirror Relationship Broken-off (ocumEvtSnapmirrorRelationshipStateBrokenoff)	Risk	SnapMirror relationship	Error
SnapMirror Relationship Initialize Failed (ocumEvtSnapmirrorRelationshipInitializeFailed)	Risk	SnapMirror relationship	Error
SnapMirror Relationship Update Failed (ocumEvtSnapmirrorRelationshipUpdateFailed)	Risk	SnapMirror relationship	Error
SnapMirror Relationship Lag Error (ocumEvtSnapMirrorRelationshipLagError)	Risk	SnapMirror relationship	Error
SnapMirror Relationship Lag Warning (ocumEvtSnapMirrorRelationshipLagWarning)	Risk	SnapMirror relationship	Warning
SnapMirror Relationship Resync Failed (ocumEvtSnapmirrorRelationshipResyncFailed)	Risk	SnapMirror relationship	Error
SnapMirror Relationship Deleted ocumEvtSnapmirrorRelationshipDeleted	Risk	SnapMirror relationship	Warning

Snapshot events

Snapshot events provide information about the status of snapshots which enables you to monitor the snapshots for potential problems. The events are grouped by impact area, and include the event name, trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Snapshot Auto-delete Disabled (Not applicable)	Event	Volume	Information
Snapshot Auto-delete Enabled (Not applicable)	Event	Volume	Information
Snapshot Auto-delete Configuration Modified (Not applicable)	Event	Volume	Information

SnapVault relationship events

SnapVault relationship events provide you with information about the status of your SnapVault relationships so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: protection

Event name (Trap name)	Impact level	Source type	Severity
SnapVault Relationship Unhealthy (ocumEvtSnapVaultRelationshipUnhealthy)	Risk	SnapMirror relationship	Warning
SnapVault Relationship Broken-off (ocumEvtSnapVaultRelationshipStateBrokenoff)	Risk	SnapMirror relationship	Error
SnapVault Relationship Initialize Failed (ocumEvtSnapVaultRelationshipInitializeFailed)	Risk	SnapMirror relationship	Error
SnapVault Relationship Update Failed (ocumEvtSnapVaultRelationshipUpdateFailed)	Risk	SnapMirror relationship	Error
SnapVault Relationship Lag Error (ocumEvtSnapVaultRelationshipLagError)	Risk	SnapMirror relationship	Error
SnapVault Relationship Lag Warning (ocumEvtSnapVaultRelationshipLagWarning)	Risk	SnapMirror relationship	Warning
SnapVault Relationship Resync Failed (ocumEvtSnapvaultRelationshipResyncFailed)	Risk	SnapMirror relationship	Error

Storage failover settings events

Storage failover (SFO) settings events provide you with information about whether your storage failover is disabled or not configured so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Storage Failover Interconnect One Or More Links Down (ocumEvtSfoInterconnectOneOrMoreLinksDown)	Risk	Node	Warning
Storage Failover Disabled (ocumEvtSfoSettingsDisabled)	Risk	Node	Error
Storage Failover Not Configured (ocumEvtSfoSettingsNotConfigured)	Risk	Node	Error
Storage Failover State - Takeover (ocumEvtSfoStateTakeover)	Risk	Node	Warning
Storage Failover State - Partial Giveback (ocumEvtSfoStatePartialGiveback)	Risk	Node	Error

Event name (Trap name)	Impact level	Source type	Severity
Storage Failover Node Status Down (ocumEvtSfoNodeStatusDown)	Risk	Node	Error
Storage Failover Takeover Not Possible (ocumEvtSfoTakeoverNotPossible)	Risk	Node	Error

Storage services events

Storage services events provide you with information about the creation and subscription of storage services so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Storage Service Created (Not applicable)	Event	Storage service	Information
Storage Service Subscribed (Not applicable)	Event	Storage service	Information
Storage Service Unsubscribed (Not applicable)	Event	Storage service	Information

Impact area: protection

Event name (Trap name)	Impact level	Source type	Severity
Unexpected Deletion of Managed SnapMirror Relationship ocumEvtStorageServiceUnsupportedRelationshipDeletion	Risk	Storage service	Warning
Unexpected Deletion of Storage Service Member Volume (ocumEvtStorageServiceUnexpectedVolumeDeletion)	Incident	Storage service	Critical

Storage shelf events

Storage shelf events tell you if your storage shelf has abnormal so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Abnormal Voltage Range (ocumEvtShelfVoltageAbnormal)	Risk	Storage shelf	Warning

Event name (Trap name)	Impact level	Source type	Severity
Abnormal Current Range (ocumEvtShelfCurrentAbnormal)	Risk	Storage shelf	Warning
Abnormal Temperature (ocumEvtShelfTemperatureAbnormal)	Risk	Storage shelf	Warning

User and group quota events

User and group quota events provide you with information about the capacity of the user and user group quota as well as the file and disk limits so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
User or Group Quota Disk Space Soft Limit Breached (ocumEvtUserOrGroupQuotaDiskSpaceSoftLimitBreached)	Risk	User or group quota	Warning
User or Group Quota Disk Space Hard Limit Reached (ocumEvtUserOrGroupQuotaDiskSpaceHardLimitReached)	Incident	User or group quota	Critical
User or Group Quota File Count Soft Limit Breached (ocumEvtUserOrGroupQuotaFileCountSoftLimitBreached)	Risk	User or group quota	Warning
User or Group Quota File Count Hard Limit Reached (ocumEvtUserOrGroupQuotaFileCountHardLimitReached)	Incident	User or group quota	Critical

Volume events

Volume events provide information about the status of volumes which enables you to monitor for potential problems. The events are grouped by impact area, and include the event name, trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Volume Restricted (ocumEvtVolumeRestricted)	Risk	Volume	Warning
Volume Offline (ocumEvtVolumeOffline)	Incident	Volume	Critical
Volume Unmounted (Not applicable)	Event	Volume	Information

Event name (Trap name)	Impact level	Source type	Severity
Volume Mounted (Not applicable)	Event	Volume	Information
Volume Remounted (Not applicable)	Event	Volume	Information
Volume Junction Path Inactive (ocumEvtVolumeJunctionPathInactive)	Risk	Volume	Warning
Volume Autosize Enabled (Not applicable)	Event	Volume	Information
Volume Autosize-Disabled (Not applicable)	Event	Volume	Information
Volume Autosize Maximum Capacity Modified (Not applicable)	Event	Volume	Information
Volume Autosize Increment Size Modified (Not applicable)	Event	Volume	Information

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Thin-Provisioned Volume Space At Risk (ocumThinProvisionVolumeSpaceAtRisk)	Risk	Volume	Warning
Volume Space Full (ocumEvtVolumeFull)	Risk	Volume	Error
Volume Space Nearly Full (ocumEvtVolumeNearlyFull)	Risk	Volume	Warning
Volume Snapshot Reserve Space Full (ocumEvtSnapshotFull)	Risk	Volume	Warning
Too Many Snapshot Copies (ocumEvtSnapshotTooMany)	Risk	Volume	Error
Volume Qtree Quota Overcommitted (ocumEvtVolumeQtreeQuotaOvercommitted)	Risk	Volume	Error
Volume Qtree Quota Nearly Overcommitted (ocumEvtVolumeQtreeQuotaAlmostOvercommitted)	Risk	Volume	Warning
Volume Growth Rate Abnormal (ocumEvtVolumeGrowthRateAbnormal)	Risk	Volume	Warning
Volume Days Until Full (ocumEvtVolumeDaysUntilFullSoon)	Risk	Volume	Error
Volume Space Guarantee Disabled (Not applicable)	Event	Volume	Information

Event name (Trap name)	Impact level	Source type	Severity
Volume Space Guarantee Enabled (Not Applicable)	Event	Volume	Information
Volume Space Guarantee Modified (Not applicable)	Event	Volume	Information
Volume Snapshot Reserve Days Until Full (ocumEvtVolumeSnapshotReserveDaysUntilFullSoon)	Risk	Volume	Error

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Volume Renamed (Not applicable)	Event	Volume	Information
Volume Discovered (Not applicable)	Event	Volume	Information
Volume Deleted (Not applicable)	Event	Volume	Information

Volume move status events

Volume move status events tell you about the status of your volume move so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Volume Move Status: In Progress (Not applicable)	Event	Volume	Information
Volume Move Status - Failed (ocumEvtVolumeMoveFailed)	Risk	Volume	Error
Volume Move Status: Completed (Not applicable)	Event	Volume	Information
Volume Move - Cutover Deferred (ocumEvtVolumeMoveCutoverDeferred)	Risk	Volume	Warning

SVM events

Storage Virtual Machine (SVM) events provide you with information about the status of your SVMs so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
SVM CIFS Service Down (ocumEvtVserverCifsServiceStatusDown)	Incident	SVM	Critical
SVM FC/FCoE Service Down (ocumEvtVserverFcServiceStatusDown)	Incident	SVM	Critical
SVM iSCSI Service Down (ocumEvtVserverIscsiServiceStatusDown)	Incident	SVM	Critical
SVM NFS Service Down (ocumEvtVserverNfsServiceStatusDown)	Incident	SVM	Critical
SVM CIFS Service Not Configured (Not applicable)	Event	SVM	Information
SVM FC/FCoE Service Not Configured (Not applicable)	Event	SVM	Information
SVM iSCSI Service Not Configured (Not applicable)	Event	SVM	Information
SVM NFS Service Not Configured (Not applicable)	Event	SVM	Information
SVM Stopped (ocumEvtVserverDown)	Risk	SVM	Warning
SVM with Infinite Volume Storage Not Available (ocumEvtVserverStorageNotAvailable)	Incident	SVM SVMs with Infinite Volume	Critical
SVM with Infinite Volume Storage Partially Available (ocumEvtVserverStoragePartiallyAvailable)	Risk	SVM SVMs with Infinite Volume	Error
SVM with Infinite Volume Namespace Mirror Constituents Having Availability Issues (ocumEvtVserverNsMirrorAvailabilityHavingIssues)	Risk	SVM SVMs with Infinite Volume	Warning

Impact area: capacity

The following capacity events apply only to SVMs with Infinite Volume.

Event name (Trap name)	Impact level	Source type	Severity
SVM with Infinite Volume Space Full (ocumEvtVserverFull)	Risk	SVM	Error
SVM with Infinite Volume Space Nearly Full (ocumEvtVserverNearlyFull)	Risk	SVM	Warning
SVM with Infinite Volume Snapshot Usage Limit Exceeded (ocumEvtVserverSnapshotUsageExceeded)	Risk	SVM	Warning
SVM with Infinite Volume Namespace Space Full (ocumEvtVserverNamespaceFull)	Risk	SVM	Error
SVM with Infinite Volume Namespace Space Nearly Full (ocumEvtVserverNamespaceNearlyFull)	Risk	SVM	Warning

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
SVM Discovered (Not applicable)	Event	SVM	Information
SVM Deleted (Not applicable)	Event	Cluster	Information
SVM Renamed (Not applicable)	Event	SVM	Information

SVM storage class events

Storage Virtual Machine (SVM) storage class events provide you with information about the status of your storage classes so that you can monitor for potential problems. SVM storage classes exist only in SVMs with Infinite Volume. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

The following SVM storage class events apply only to SVMs with Infinite Volume.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
SVM Storage Class Not Available (ocumEvtVserverStorageClassNotAvailable)	Incident	Storage class	Critical
SVM Storage Class Partially Available (ocumEvtVserverStorageClassPartiallyAvailable)	Risk	Storage class	Error

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
SVM Storage Class Space Nearly Full (ocumEvtVserverStorageClassNearlyFull)	Risk	Storage class	Warning
SVM Storage Class Space Full (ocumEvtVserverStorageClassFull)	Risk	Storage class	Error
SVM Storage Class Snapshot Usage Limit Exceeded (ocumEvtVserverStorageClassSnapshotUsageExceeded)	Risk	Storage class	Warning

What alerts are

While events occur continuously, the OnCommand Unified Manager server generates an alert only when an event meets specified filter criteria. You can choose the events for which alerts should be generated—for example, when a space threshold is exceeded or an object goes offline.

Filter criteria include object class, name, or event severity.

Related concepts

What events are on page 73

Adding an alert

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

Before you begin

- You must have configured notification settings such as the user email address, SMTP server, and SNMP trap host so that the Unified Manager server can use these settings to send notifications to users when an event is generated.
- You must know the resources and events for which you want to trigger the alert, and the user names or email addresses of the users that you want to notify.
- You must have added scripts to Unified Manager by using the Manage Scripts page.
- You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can create an alert based on resources, events, or both.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Alerts**.
3. In the **Manage Alerts** page, click **Add**.

4. In the **Add Alert** dialog box, perform the following steps:
    a. Click **Name**, and enter a name and description for the alert.
    b. Click **Resources**, and select the resources to be included in or excluded from the alert.

        You can set a filter by specifying a text string in the **Name contains** field to select a group of resources. Based on the text string that you specify, the list of available resources displays only those resources that match the filter rule. The text string that you specify is case-sensitive.

        If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.
    c. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.
    d. Click **Actions**, and select the users that you want to notify, choose the notification frequency, and assign a script to be executed when an alert is generated.

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

        You can also choose to notify users through SNMP traps.
5. Click **Save**.

Example of adding an alert

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

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

Perform the following steps in the Add Alert dialog box:

1. Click **Name**, and enter `Test` in the **Alert Name** field.
2. Click **Resources**, and in the Include tab, select **Volumes** from the drop-down list.
    a. Enter **`abc`** in the **Name contains** field to display the volumes whose name contains “abc”.
    b. Select **<<All Volumes whose name contains 'abc'>>** from the Available Resources area, and move it to the Selected Resources area.
    c. Click **Exclude**, and enter **`xyz`** in the **Name contains** field, and then click **Add**.
3. Click **Events**, and select **Critical** from the Event Severity field.
4. Select **All Critical Events** from the Matching Events area, and move it to the Selected Events area.

5. Click **Actions**, and enter **`sample@domain.com`** in the Alert these users field.
6. Select **Remind every 15 minutes** to notify the user every 15 minutes.
   You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.
7. In the Select Script to Execute menu, select **Test** script .
8. Click **Save**.

Related concepts

What alerts are on page 102
Guidelines for adding alerts on page 104

Related tasks

Adding a user on page 379
Configuring notification settings on page 75

Guidelines for adding alerts

You can add alerts based on a resource, such as a cluster, node, aggregate, or volume, and events of a particular severity type. As a best practice, you can add an alert for any of your critical objects after you have added the cluster to which the object belongs.

You can use the following guidelines and considerations to create alerts to manage your systems effectively:

- Alert description
  You should provide a description for the alert so that it helps you track your alerts effectively.
- Resources
  You should decide which physical or logical resource requires an alert. You can include and exclude resources, as required. For example, if you want to closely monitor your aggregates by configuring an alert, you must select the required aggregates from the list of resources.
- Event severity
  You should decide if an event of a specified severity type (Critical, Error, Warning) should trigger the alert and, if so, which severity type.
- Event name
  If you add an alert based on the type of event generated, you should decide which events require an alert
- Actions
  You must provide the user names and email addresses of the users who receive the notification. You can also specify an SNMP trap as a mode of notification. You can associate your scripts to an alert so that they are executed when an alert is generated.
- Notification frequency
  You can configure an alert to repeatedly send notification to the recipients for a specified time. You should determine the time from which the event notification is active for the alert. If you want the event notification to be repeated until the event is acknowledged, you should determine how often you want the notification to be repeated.
- Execute Script
  You can associate your script with an alert. Your script is executed when the alert is generated.

Testing an alert

You can test an alert to verify that you have configured it correctly. When an event is triggered, an alert is generated, and an alert email is sent to the configured recipients. You can verify whether the notification is sent and whether your script is executed by using the test alert.

Before you begin

- You must have configured notification settings such as the email address of the recipients, SMTP server, and SNMP trap.
  The Unified Manager server can use these settings to send notifications to users when an event is generated.
- You must have assigned a script and configured the script to run when the alert is generated.
- You must have the OnCommand Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Alerts**.
3. In the **Manage Alerts** page, select the alert that you want to test, and then click **Test**.

   A test alert email is sent to the email addresses that you specified while creating the alert.

Related tasks

Adding a user on page 379
Configuring notification settings on page 75

Viewing alerts

You can view the list of alerts that is created for various events from the Manage Alerts page. You can also view alert properties such as the alert description, notification method and frequency, events that trigger the alert, email recipients of the alerts, and affected resources such as clusters, aggregates, and volumes.

Before you begin

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

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Alerts**.

   The list of alerts is displayed in the Manage Alerts page.

Related tasks

Adding a user on page 379

Editing an alert

You can edit alert properties such as the resource with which the alert is associated, events, recipients, notification options, notification frequency, and associated scripts.

Before you begin

You must have the OnCommand Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Alerts**.
3. In the **Manage Alerts** page, select the alert that you want to edit, and click **Edit**.
4. In the **Edit Alert** dialog box, edit the name, resources, events, and actions sections, as required.

   You can change or remove the script that is associated with the alert.
5. Click **Save**.

Related concepts

What alerts are on page 102

Related tasks

Adding a user on page 379

Deleting alerts

You can delete an alert when it is no longer required. For example, you can delete an alert that was created for a particular resource when the resource is no longer monitored by Unified Manager.

Before you begin

You must have the OnCommand Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Alerts**.
3. On the **Manage Alerts** page, select the alerts that you want to delete, and click **Delete**.
4. Click **Yes** to confirm the delete request.

Related tasks

Adding a user on page 379

Description of Event windows and dialog boxes

Events notify you about any issues in your environment. You can use the Events page and Event details page to monitor all the events. You can use the Notification Setup Options dialog box to configure notification. You can use the Manage Events page to disable or enable events.

Event Retention Settings

You can configure the event settings to automatically delete events (information, resolved, or obsolete) after a specified time and at a specified frequency. You can also delete these events manually through the Event Retention Settings.

You must have the OnCommand Administrator or Storage Administrator role.

Event Settings

You can configure the following options:

Delete Information, Resolved, and Obsolete Events Older Than

Enables you to specify the retention period after which information, resolved, and obsolete events are removed from the management server.

The default value is 180 days. Retaining the events for more than 180 days affects the performance and is not recommended. The lower limit for the event retention period is 7 days, although there is no upper limit.

Delete Schedule

Enables you to specify the frequency at which all the information, resolved, and obsolete events that have exceeded their age limit are automatically deleted from the management server. The possible values are Daily, Weekly, or Monthly.

The default value is Daily.

Delete Now

Enables you to manually delete all the information, resolved, and obsolete events that have exceeded their specified retention period.

Command buttons

The command buttons enable you to save or cancel the setup options:

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Configuring event settings on page 75

Notification Setup Options dialog box

You can configure the management server to send notifications when an event is generated or when it is assigned to a user. You can also configure the notification mechanisms. For example, notifications can be sent as emails or SNMP traps.

You must have the OnCommand Administrator or Storage Administrator role.

Email

This area enables you to configure the following email settings for alert notification:

From Address

Specifies the email address from which the alert notification is sent, also specifies the from address for a report shared.

SMTP Server

This area enables you to configure the following SMTP server settings:

Host Name

Specifies the host name of your SMTP host server, which is used to send the alert notification to the specified recipients.

User Name

Specifies the SMTP user name. SMTP user name is required only when the SMTPAUTH is enabled in the SMTP server.

Password

Specifies the SMTP password. SMTP user name is required only when the SMTPAUTH is enabled in the SMTP server.

Default Port

Specifies the port that is used by the SMTP host server to send alert notification.

The default value is 25.

Use secure connection

Selecting this box provides secure communication between the SMTP server and the management server.

SNMP

This area enables you to configure the following SNMP trap settings:

Version

Specifies the SNMP version you want to use depending on the type of security you require. Options include Version 1, Version 3, Version 3 with Authentication, and Version 3 with Authentication and Encryption. The default value is Version 1.

Trap Destination Host

Specifies the host name or IP address (IPv4 or IPv6) that receives the SNMP traps that are sent by the management server.

Outbound Trap Port

Specifies the port through which the SNMP server receives the traps that are sent by the management server.

The default value is 162.

Community

Specifies the community name that is used by the SNMP server to authenticate the traps that are sent from the management server. The community name is only available with SNMP Version 1.

Engine ID

Specifies the unique identifier of the SNMP agent and is automatically generated by the management server. Engine ID is available with SNMP Version 3, SNMP Version 3 with Authentication, and SNMP Version 3 with Authentication and Encryption.

User Name

Specifies the SNMP user name. User name is available with SNMP Version 3, SNMP Version 3 with Authentication, and SNMP Version 3 with Authentication and Encryption.

Authentication Protocol

Specifies the protocol used to authenticate a user. Protocol options include MD5 and SHA. MD5 is the default value. Authentication protocol is available with SNMP Version 3 with Authentication and SNMP Version 3 with Authentication and Encryption.

Authentication Password

Specifies the password used when authenticating a user. Authentication password is available with SNMP Version 3 with Authentication and SNMP Version 3 with Authentication and Encryption.

Privacy Protocol

Specifies the privacy protocol used to encrypt SNMP messages. Protocol options include AES 128 and DES. The default value is AES 128. Privacy protocol is available with SNMP Version 3 with Authentication and Encryption.

Privacy Password

Specifies the password when using privacy protocol. Privacy password is available with SNMP Version 3 with Authentication and Encryption.

Command buttons

The command buttons enable you to save or cancel the setup options:

Restore to Factory Defaults

Enables you to restore the configuration settings to the factory default values.

Save

Saves the configuration settings for the selected option.

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Configuring notification settings on page 75

Events page

The Events page enables you to view a list of current events and their properties. You can perform tasks such as acknowledging, resolving, and assigning events. You can also add an alert to a specific event.

- *Command buttons* on page 109
- *Events list* on page 110
- *Filters pane* on page 111

Command buttons

The command buttons enable you to perform the following tasks:

Assign To

Enables you to select the user to whom the event is assigned. You can select one of the following:

Me

Assigns the event to you.

When you assign an event to yourself, your user name and the time when you assigned the event is added in the events list for the selected events.

Another user

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

When you assign an event to a user, the user's name and the time when the event was assigned are added in the events list for the selected events. You can also unassign events by leaving the ownership field blank.

Acknowledge

Acknowledges the selected events.

When you acknowledge an event, your user name and the time when you acknowledged the event are added in the events list for the selected events. When you acknowledge an event, you are responsible for managing that event.

Mark As Resolved

Enables you to change the event state to resolved.

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

Add Alert

Displays the Add Alert dialog box, which enables you to add alerts for the selected events.

Export

Enables you to export details of all new and acknowledged events of severity type Critical, Error, and Warning to a comma-separated values (`.csv`) file.

Events list

Displays details of all the events that occurred, based on the recent timestamp. By default, events of severity type Critical, Error, and Warning and events of state New and Acknowledged are displayed.

Triggered Time

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp when the event was generated is displayed. You can use the time period filter to narrow your search.

Status

Displays the severity of the event. You can filter this column to display events of a specific severity type or types. The event severity types are Critical (✖), Error (❗), Warning (⚠), and Information (ℹ).

State

Displays the event state: New, Acknowledged, Resolved, or Obsolete. You can filter this column to show events of a specific state.

Impact Level

Displays whether the event is categorized as an incident, risk, or an informational event.

Impact Area

Displays whether the event is a capacity, availability, performance, protection, or configuration related event.

Name

Displays the event names. You can select an event to display the event details.

Source

Displays the name of the object where the event has occurred.

Source Type

Displays the object type (for example, Storage Virtual Machine (SVM), volume, or qtree) with which the event is associated.

Assigned To

Displays the name of the user to whom the event is assigned.

Assigned Time

Displays the time that has elapsed since the event was assigned to a user. If the time elapsed exceeds a week, the timestamp when the event was assigned to a user is displayed. By default, this column is hidden.

Notes

Displays the number of notes that are added for an event.

Acknowledged By

Displays the name of the user who acknowledged the event. The field is blank if the event is not acknowledged. By default, this column is hidden.

Acknowledged Time

Displays the time that has elapsed since the event was acknowledged. If the time elapsed exceeds a week, the timestamp when the event was acknowledged is displayed. By default, this column is hidden.

Resolved By

Displays the name of the user who resolved the event. The field is blank if the event is not resolved. By default, this column is hidden.

Resolved Time

Displays the time that has elapsed since the event was resolved. If the time elapsed exceeds a week, the timestamp when the event was resolved is displayed. By default, this column is hidden.

Obsoleted Time

Displays the time when the state of the event became Obsolete. By default, this column is hidden.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the events list. You can select filters related to the Status, State, Impact Level, Impact Area, Source Type, Assigned To, and Annotation.

Note: The filters specified in the Filters pane override the filters specified for the columns in the events list.

Related tasks

Viewing event details on page 77
Assigning events on page 78
Viewing unassigned events on page 79
Acknowledging and resolving events on page 80

Adding an alert on page 102
Exporting data to CSV files on page 31

Event details page

From the Event details page, you can view the details of a selected event, such as the event severity, impact level, impact area, and event source. You can also view additional information about the selected event in the Notes and Updates area section, which is provided by the user who previously worked on that event.

- *Command buttons* on page 112
- *Summary area* on page 112
- *Notes and Updates area* on page 114
- *Scope of Impact* on page 114
- *Possible Causes* on page 114
- *Possible Effect* on page 114
- *Resources that Might be Impacted* on page 114
- *Suggested Corrective Actions area* on page 114

Command buttons

The command buttons enable you to perform the following tasks:

Assign To

Me

Assigns the event to you.

Another user

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

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

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

Acknowledge

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

Mark As Resolved

Enables you to change the event state to Resolved.

Add Alert

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

View Events

Navigates to the Events page.

Summary area

You can view the following event details:

Severity

Displays the severity of the event.

The event severity types are Critical (✖), Error (❗), Warning (⚠), and Information (ℹ).

State

Displays the event state: New, Acknowledged, Resolved, or Obsolete.

Impact Level

Displays whether the event is categorized as an incident, risk, or an informational event.

Impact Area

Displays whether the event is a capacity, availability, protection, performance, or configuration related event.

Obsoleted Cause

Displays the reason the event is now obsolete.

Source

Displays the full name of the object, along with the type of object with which the event is associated.

The value is displayed as “Unknown” when Data ONTAP does not provide a valid user name because of SecD errors.

Source Type

Displays the type of storage object for which the event is created.

Source Annotations

Displays the annotation pairs associated with the impacted object.

Source Groups

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

Source Annotations

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

Source Type

Displays the object type (for example, Storage Virtual Machine (SVM), volume, qtree, or root or data aggregate) with which the event is associated.

Acknowledged By

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

This field is blank if the event is not acknowledged.

Resolved By

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

This field is blank if the event is not resolved.

Assigned To

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

Triggered Time

Displays the time that has elapsed since the event was generated.

If the time elapsed exceeds a week, the timestamp when the event was generated is displayed.

Trigger Condition

Displays information about the cause of the event.

Alert Settings

The following information about alerts is displayed:

- If there are no alerts associated with the selected event, an **Add** link is displayed.
  You can open the Add Alert dialog box by clicking the link.
- If there is one alert associated with the selected event, the alert name is displayed.
  You can open the Edit Alert dialog box by clicking the link.
- If there is more than one alert associated with the selected event, the number of alerts is displayed.
  You can open the Alerts page by clicking the link to view more details about these alerts.

Alerts that are disabled are not displayed.

Notes and Updates area

Displays information that was added by the user who last addressed the generated event, based on the recent timestamp. You can also view the time when the information was added.

Post

Enables you to display the information that you added.

Scope of Impact area

Graphically displays the resources that are impacted because of the generated event. You can click the name or count link for each of the resources to view more details of the impacted resources.

This area is displayed only for some events, such as the Some Failed Disks, Aggregate Degraded, MetroCluster Aggregate Mirroring Degraded, Node To FC Switch All FC-Initiator Links Down, and All Links Between MetroCluster Partners Down.

Possible Causes area

Displays one or more causes that generated the event. You can click the name of the resource to view more details about that resource.

This area is displayed only for some events, such as Aggregate Degraded, Cluster Lacks Spare Disks, MetroCluster Spare Disks Left Behind, and All Links Between MetroCluster Partners Down.

Possible Effect

Displays the effect of the generated event.

This area is displayed only for MetroCluster events, such as All Links Between MetroCluster Partners Down, MetroCluster Aggregate Mirroring Degraded, and Node To FC Switch All FC-Initiator Links Down.

Resources that Might Be Impacted area

Displays the resources that might be impacted because of the generated event. You can click the name of the resource to view more details about the impacted resource.

This area is displayed only for some events, such as Aggregate Degraded, and Some Failed Disks.

Suggested Corrective Action area

Displays the actions that you can perform to address the issues.

This area is displayed only for some events, such as Volume Space Nearly Full, Volume Space Full, Aggregate Degraded, Some Failed Disks, Node To FC Switch All FC-Initiator Links Down, and All Links Between MetroCluster Partners Down.

Related tasks

Viewing event details on page 77
Assigning events on page 78
Acknowledging and resolving events on page 80
Adding an alert on page 102
Adding and reviewing notes about an event on page 79

Manage Events page

The Manage Events page displays the list of events that are disabled, and provides information such as the associated object type and severity of the event. You can also perform tasks such as disabling or enabling events globally.

You can access this page only if you have the OnCommand Administrator or Storage Administrator role.

- *Command buttons* on page 115
- *List view* on page 115

Command buttons

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

Disable

Launches the Disable Events dialog box, which you can use to disable events.

Enable

Enables selected events that you had chosen to disable previously.

Subscribe to EMS Events

Launches the Subscribe to EMS Events dialog box, which enables you to subscribe to receive specific Event Management System (EMS) events from the clusters that you are monitoring. The EMS collects information about events that occur on the cluster. When a notification is received for a subscribed EMS event, a Unified Manager event is generated with the appropriate severity.

Event Retention Settings

Launches the Event Retention Settings dialog box, which enables you to specify the retention period after which the information, resolved, and obsolete events are removed from the management server. The default retention value is 180 days.

List view

The List view displays (in tabular format) information about events that are disabled. You can use the column filters to customize the data that is displayed.

Event

Displays the name of the event that is disabled.

Severity

Displays the severity of the event. The severity can be Critical, Error, Warning, or Information.

Source Type

Displays the source type for which the event is generated.

Related tasks

Disabling or enabling events on page 77

Disable Events dialog box

The Disable Events dialog box displays the list of event types for which you can disable events. You can disable events for an event type based on a particular severity or for a set of events.

You must have the OnCommand Administrator or Storage Administrator role.

- *Event properties area* on page 116
- *Command buttons* on page 116

Event Properties area

The Event Properties area specifies the following event properties:

Event Severity

Enables you to select events based on the severity type, which can be Critical, Error, Warning, or Information.

Event Name Contains

Enables you to filter events whose name contains the specified characters.

Matching events

Displays the list of events matching the event severity type and the text string you specify.

Disable events

Displays the list of events that you have selected for disabling.

The severity of the event is also displayed along with the event name.

Command buttons

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

Save and close

Disables the event type and closes the Disable Events dialog box.

Cancel

Discards the changes and closes the Disable Events dialog box.

Related tasks

Disabling or enabling events on page 77

Description of Alert windows and dialog boxes

You should configure alerts to receive notifications about events by using the Add Alert dialog box. You can also view the list of alerts from the Manage Alerts page.

Manage Alerts page

The Manage Alerts page displays a list of alerts and provides information about the alert name, status, notification method, and notification frequency. You can also add, edit, remove, enable, or disable alerts from this page.

You must have the OnCommand Administrator or Storage Administrator role.

- *Command buttons* on page 117
- *List view* on page 117
- *Details area* on page 117

Command buttons

Add

Displays the Add Alert dialog box, which enables you to add new alerts.

Edit

Displays the Edit Alert dialog box, which enables you to edit selected alerts.

Delete

Deletes the selected alerts.

Enable

Enables the selected alerts to send notifications.

Disable

Disables the selected alerts when you want to temporarily stop sending notifications.

Test

Tests the selected alerts to verify their configuration after being added or edited.

List view

The list view displays, in tabular format, information about the alerts that are created. You can use the column filters to customize the data that is displayed. You can also select an alert to view more information about it in the details area.

Status

Specifies whether an alert is enabled (✓) or disabled (⊘).

Alert

Displays the name of the alert.

Description

Displays a description for the alert.

Notification Method

Displays the notification method that is selected for the alert. You can notify users through email or SNMP traps.

Notification Frequency

Specifies the frequency (in minutes) with which the management server continues to send notifications until the event is acknowledged, resolved, or moved to the Obsolete state.

Details area

The details area provides more information about the selected alert.

Alert Name

Displays the name of the alert.

Alert Description

Displays a description for the alert.

Events

Displays the events for which you want to trigger the alert.

Resources

Displays the resources for which you want to trigger the alert.

Includes

Displays the group of resources for which you want to trigger the alert.

Excludes

Displays the group of resources for which you do not want to trigger the alert.

Notification Method

Displays the notification method for the alert.

Notification Frequency

Displays the frequency with which the management server continues to send alert notifications until the event is acknowledged, resolved, or moved to the Obsolete state.

Script Name

Displays the name of the script associated with the selected alert. This script is executed when an alert is generated.

Email Recipients

Displays the email addresses of users who receive the alert notification.

Related tasks

Adding an alert on page 102
Editing an alert on page 106
Testing an alert on page 105
Deleting alerts on page 106
Viewing alerts on page 105

Add Alert dialog box

You can create alerts to notify you when a particular event is generated, so that you can address the issue quickly and thereby minimize impact to your environment. You can create alerts for a single resource or a set of resources, and for events of a particular severity type. You can also specify the notification method and frequency of the alerts.

You must have the OnCommand Administrator or Storage Administrator role.

Name

This area enables you to specify a name and description for the alert:

Alert Name

Enables you to specify an alert name.

Alert Description

Enables you to specify a description for the alert.

Resources

This area enables you to select an individual resource or group the resources based on a dynamic rule for which you want to trigger the alert. A *dynamic rule* is the set of resources filtered based on the text string you specify. You can search for resources by selecting a resource type from the drop-down list or you can specify the exact resource name to display a specific resource.

If you are creating an alert from any of the storage object details pages, the storage object is automatically included in the alert.

Include

Enables you to include the resources for which you want to trigger alerts. You can specify a text string to group resources that match the string and select this group to be included in the alert. For example, you can group all volumes whose name contains the “abc” string.

Exclude

Enables you to exclude resources for which you do not want to trigger alerts. For example, you can exclude all volumes whose name contains the “xyz” string.

The Exclude tab is displayed only when you select all resources of a particular resource type: for example, <<All Volumes>> or <<All Volumes whose name contains '*xyz*'>>.

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

Events

This area enables you to select the events for which you want to create the alerts. You can create alerts for events based on a particular severity or for a set of events.

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

Event Severity

Enables you to select events based on the severity type, which can be Critical, Error, or Warning.

Event Name Contains

Enables you to select events whose name contains specified characters.

Actions

This area enables you to specify the users that you want to notify when an alert is triggered. You can also specify the notification method and the frequency of notification.

Alert these users

Enables you to specify the email address or user name of the user to receive notifications.

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

Notification Frequency

Enables you to specify the frequency with which the management server sends notifications until the event is acknowledged, resolved, or moved to the obsolete state.

You can choose the following notification methods:

- Notify only once
- Notify at a specified frequency

- Notify at a specified frequency within the specified time range

Issue SNMP trap

Selecting this box enables you to specify whether SNMP traps should be sent to the globally configured SNMP host.

Execute Script

Enables you to add your custom script to the alert. This script is executed when an alert is generated.

Command buttons

Save

Creates an alert and closes the Add Alert dialog box.

Cancel

Discards the changes and closes the Add Alert dialog box.

Related tasks

Adding an alert on page 102

Edit Alert dialog box

You can edit alert properties such as the resource with which the alert is associated, events, script, and notification options.

Name

This area enables you to edit the name and description for the alert.

Alert Name

Enables you to edit the alert name.

Alert Description

Enables you to specify a description for the alert.

Alert State

Enables you to enable or disable the alert.

Resources

This area enables you to select an individual resource or group the resources based on a dynamic rule for which you want to trigger the alert. You can search for resources by selecting a resource type from the drop-down list or you can specify the exact resource name to display a specific resource.

Include

Enables you to include the resources for which you want to trigger alerts. You can specify a text string to group resources that match the string and select this group to be included in the alert. For example, you can group all volumes whose name contains the “vol0” string.

Exclude

Enables you to exclude resources for which you do not want to trigger alerts. For example, you can exclude all volumes whose name contains the “xyz” string.

Note: The Exclude tab is displayed only when you select all resources of a particular resource type—for example, <<All Volumes>> or <<All Volumes whose name contains '*xyz*'>>.

Events

This area enables you to select the events for which you want to trigger the alerts. You can trigger an alert for events based on a particular severity or for a set of events.

Event Severity

Enables you to select events based on the severity type, which can be Critical, Error, or Warning.

Event Name Contains

Enables you to select events whose name contains the specified characters.

Actions

This area enables you to specify the notification method and the frequency of notification.

Alert these users

Enables you to edit the email address or user name, or specify a new email address or user name to receive notifications.

Notification Frequency

Enables you to edit the frequency with which the management server sends notifications until the event is acknowledged, resolved, or moved to the obsolete state.

You can choose the following notification methods:

- Notify only once
- Notify at a specified frequency
- Notify at a specified frequency within the specified time range

Issue SNMP trap

Enables you to specify whether SNMP traps should be sent to the globally configured SNMP host.

Execute Script

Enables you to associate a script with the alert. This script is executed when an alert is generated.

Command buttons

Save

Saves the changes and closes the Edit Alert dialog box.

Cancel

Discards the changes and closes the Edit Alert dialog box.

Related tasks

Editing an alert on page 106

Managing scripts

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

How scripts work with alerts

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

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

The script uses the following arguments for execution:

- *`-eventID`*
- *`-eventSourceID`*
- *`-eventSourceName`*
- *`-eventSourceType`*
- *`-eventState`*
- *`-eventArgs`*

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

Example for obtaining arguments from scripts

```
print "$ARGV[0] : $ARGV[1]\n"
print "$ARGV[2] : $ARGV[3]\n"
```

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

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

Related references

Manage Scripts page on page 124

Adding scripts

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

Before you begin

- You must have created and saved the scripts that you want to add to Unified Manager.

  Note: The supported file formats for scripts are Perl, Shell, PowerShell, and `.bat` files.

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

About this task

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

Steps

1. Click **Health** > **Administration** > **Manage Scripts**.
2. In the **Manage Scripts** dialog box, click **Add**.
3. In the **Add Script** dialog box, click **Choose File** to select your script.
4. Enter a description for the script that you select.
5. Click **Add**.

Related tasks

Testing script execution on page 124

Related references

Manage Scripts page on page 124
Add Script dialog box on page 125

Deleting scripts

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

Before you begin

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

Steps

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

Related references

Manage Scripts page on page 124

Testing script execution

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

Before you begin

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

Steps

1. Click **Health** > **Administration** > **Manage Scripts**.
2. In the **Manage Scripts** page, add your test script.
3. In the **Manage Alerts** page, perform one of the following actions:

To...	Do this...
Add an alert	**a.** In the Manage Alerts page, click **Add**. **b.** In the Actions section, associate the alert with your test script.
Edit an alert	**a.** In the Manage Alerts page, select an alert, and then click **Edit**. **b.** In the Actions section, associate the alert with your test script.

4. Click **Save**.
5. In the **Manage Alerts** page, select the alert that you added or modified, and then click **Test**.

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

Related tasks

Adding scripts on page 123

Description of script windows and dialog boxes

The Manage Scripts menu enables you to add scripts to Unified Manager.

Manage Scripts page

The Manage Scripts page enables you to add your custom scripts to Unified Manager. You can associate these scripts with alerts to enable automatic reconfiguration of storage objects.

The Manage Scripts page includes tabs that enable you to add or delete scripts from Unified Manager.

Command buttons

Add

Displays the Add Script dialog box, which enables you to add scripts.

Delete

Deletes the selected script.

List view

The list view displays, in tabular format, the scripts that you added to Unified Manager.

Name

Displays the name of the script.

Description

Displays the description of the script.

Related concepts

How scripts work with alerts on page 122

Related tasks

Adding scripts on page 123
Deleting scripts on page 123

Add Script dialog box

The Add Script dialog box enables you to add scripts to Unified Manager. You can configure alerts with your scripts to automatically resolve events that are generated for storage objects.

You must have the OnCommand Administrator or Storage Administrator role.

Select Script File

Enables you to select a script for the alert.

Description

Enables you to specify a description for the script.

Related tasks

Adding scripts on page 123

Managing thresholds

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

What storage capacity thresholds are

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

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

Configuring global aggregate threshold values

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

- Most of the threshold options have default values for better monitoring. However, you can change the values to suit the requirements of your environment.
- The threshold values are not applicable to the root aggregate of the node.

Steps

1. Click ▦ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Thresholds** > **Aggregates**.
4. Configure the appropriate threshold values for capacity, growth, and Snapshot copies.
5. Click **Save and Close**.

Related concepts

Understanding capacity events and thresholds for node root aggregates on page 145

Related tasks

Adding a user on page 379

Configuring global volume threshold values

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. Click ▦ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Thresholds** > **Volumes**.
4. Configure the appropriate threshold values for capacity, Snapshot copies, quotas, volume growth, and inodes.
5. Click **Save and Close**.

Related tasks

Adding a user on page 379

Editing lag threshold settings for unmanaged protection relationships

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. Click ▦ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Thresholds** > **Relationships**.

4. In the **Lag** area of the **Lag Thresholds for Unmanaged Relationships** dialog box, increase or decrease the global default lag warning or error lag time percentage as required.
5. Click **Save and Close**.

Related tasks

Adding a user on page 379

Editing global threshold settings

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

About this task

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

Global threshold settings are accessible from the Setup Options dialog box. You can also modify threshold settings for individual objects, from the inventory page or the details page for that object.

Choices

- *Configuring global aggregate threshold values* on page 126

  You can edit the threshold settings for capacity, growth, and Snapshot copies for all aggregates to track any threshold breach.
- *Configuring global volume threshold values* on page 127

  You can edit the threshold settings for capacity, Snapshot copies, quotas, volume growth, overwrite reserve space, and inodes for all volumes to track any threshold breach.
- *Editing unmanaged relationship lag thresholds* on page 127

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

Editing aggregate threshold settings

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

- Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.
- The threshold values are not applicable to the root aggregate of the node.

Steps

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

Related concepts

Understanding capacity events and thresholds for node root aggregates on page 145

Related tasks

Adding a user on page 379

Editing volume threshold settings

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

Related tasks

Adding a user on page 379

Description of thresholds dialog boxes

You can use the appropriate Setup Options dialog box to configure global threshold values for aggregates and volumes.

Aggregate Thresholds Setup Options dialog box

The Aggregate Thresholds Setup Options dialog box enables you to configure global threshold values for monitored aggregates. You can set thresholds for individual aggregates or for all the aggregates globally. When you configure the options globally, the default values of the objects are modified. However, if the default values have been changed at the object level, the global values are not modified.

You must have the OnCommand Administrator or Storage Administrator role.

Events are generated when a threshold is breached. You can take corrective actions for such events.

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

You can set thresholds for the following: capacity, aggregate growth, and aggregate Snapshot copies.

- *Capacity* on page 130
- *Growth* on page 131
- *Snapshot Copies* on page 131
- *Command buttons* on page 131

Capacity area

The Capacity area enables you to set the following aggregate capacity threshold conditions:

Space Nearly Full

Specifies the percentage at which an aggregate is considered to be nearly full:

- Default value: 80 percent
  The value for this threshold must be lower than the value for the Aggregate Full threshold for the management server to generate an event.
- Event generated: Aggregate Nearly Full
- Event severity: Warning

Space Full

Specifies the percentage at which an aggregate is considered full:

- Default value: 90 percent
- Event generated: Aggregate Full
- Event severity: Error

Nearly Overcommitted

Specifies the percentage at which an aggregate is considered to be nearly overcommitted:

- Default value: 95 percent
  The value for this threshold must be lower than the value for the Aggregate Overcommitted Full threshold for the management server to generate an event.
- Event generated: Aggregate Nearly Overcommitted
- Event severity: Warning

Overcommitted

Specifies the percentage at which an aggregate is considered to be overcommitted:

- Default value: 100 percent

- Event generated: Aggregate Overcommitted
- Event severity: Error

Days Until Full

Specifies the number of days remaining before the aggregate reaches full capacity:

- Default value: 7
- Event generated: Aggregate Days Until Full
- Event severity: Error

Growth area

The Growth area enables you to set the following threshold conditions for aggregate growth:

Growth Rate

Specifies the percentage at which an aggregate's growth rate is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event:

- Default value: 1 percent
- Event generated: Aggregate Growth Rate Abnormal
- Event severity: Warning

Growth Rate Sensitivity

Specifies the factor that is applied to the standard deviation of an aggregate's growth rate. If the growth rate exceeds the factored standard deviation, an Aggregate Growth Rate Abnormal event is generated.

A lower value for growth rate sensitivity indicates that the aggregate is highly sensitive to changes in the growth rate. The range for the growth rate sensitivity is 1 through 5.

- Default value: 2

Note: If you modify the growth rate sensitivity for aggregates at the global threshold level, the change is also applied to the growth rate sensitivity for volumes at the global threshold level.

Snapshot copies area

The Snapshot copies area enables you to set the following Snapshot reserve threshold conditions:

Snapshot Reserve Full

Specifies the percentage at which an aggregate has consumed all the space reserved for Snapshot copies:

- Default value: 90 percent
- Event generated: Aggregate Snapshot Reserve Full
- Event severity: Warning

Command buttons

The command buttons enable you to save or cancel the setup options:

Restore to Factory Defaults

Enables you to restore the configuration settings to the factory default values.

Save

Saves the configuration settings for the selected option.

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related concepts

Understanding capacity events and thresholds for node root aggregates on page 145

Related tasks

Configuring global aggregate threshold values on page 126

Volume Thresholds Setup Options dialog box

The Volume Thresholds Setup Options dialog box enables you to configure global threshold values for monitored volumes. You can set thresholds for individual volumes or for all the volumes globally. When you configure the options globally, the default values of the objects are modified. However, if the default values have been changed at the object level, the global values are not modified.

You must have the OnCommand Administrator or Storage Administrator role.

Events are generated when a threshold is breached. You can take corrective actions for such events.

You can set thresholds for the following: capacity, volume Snapshot copies, quotas, volume growth, and inodes.

- *Capacity* on page 132
- *Snapshot Copies* on page 133
- *Qtree Quota* on page 133
- *Growth* on page 134
- *Inodes* on page 134
- *Command buttons* on page 135

Capacity area

The Capacity area enables you to set the following volume capacity threshold conditions:

Space Nearly Full

Specifies the percentage at which a volume is considered to be nearly full:

- Default value: 80 percent
  The value for this threshold must be lower than the value for the Volume Full threshold in order for the management server to generate an event.
- Event generated: Volume Nearly Full
- Event severity: Warning

Space Full

Specifies the percentage at which a volume is considered full:

- Default value: 90 percent
- Event generated: Volume Full

- Event Severity: Error

Days Until Full

Specifies the number of days remaining before the volume reaches full capacity:

- Default value: 7
- Event generated: Volume Days Until Full
- Event severity: Error

Snapshot copies area

The Snapshot copies area enables you to set the following threshold conditions for the Snapshot copies in the volume:

Snapshot Reserve Full

Specifies the percentage at which the space reserved for Snapshot copies is considered full:

- Default value: 90 percent
- Event generated: Volume Snapshot Reserve Full
- Event severity: Error

Days Until Full

Specifies the number of days remaining before the space reserved for Snapshot copies reaches full capacity:

- Default value: 7
- Event generated: Volume Snapshot Reserve Days Until Full
- Event severity: Error

Count

Specifies the number of Snapshot copies that can be created on a volume before the system generates the Too Many Snapshot Copies event:

- Default value: 250
- Event generated: Too Many Snapshot Copies
- Event severity: Error

Note: This field is applicable only for volumes in a cluster running Data ONTAP 8.2 or later.

Qtree Quota area

The Qtree Quota area enables you to set the following volume quota threshold conditions:

Nearly Overcommitted

Specifies the percentage at which a volume is considered to be nearly overcommitted by qtree quotas:

- Default value: 95 percent
- Event generated: Volume Qtree Quota Nearly Overcommitted
- Event severity: Warning

Overcommitted

Specifies the percentage at which a volume is considered to be overcommitted by qtree quotas:

- Default value: 100 percent
- Event generated: Volume Qtree Quota Overcommitted
- Event severity: Error

Growth

The Growth area enables you to set the following threshold conditions for volume growth:

Growth Rate

Specifies the percentage at which a volume's growth rate is considered to be normal before the system generates a Volume Growth Rate Abnormal event:

- Default value: 1 percent
- Event generated: Volume Growth Rate Abnormal
- Event severity: Warning

Growth Rate Sensitivity

Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated.

A lower value for growth rate sensitivity indicates that the volume is highly sensitive to changes in the growth rate. The range for the growth rate sensitivity is 1 through 5.

- Default value: 2

Note: If you modify the growth rate sensitivity for volumes at the global threshold level, the change is also applied to the growth rate sensitivity for aggregates at the global threshold level.

Inodes

The Inodes area enables you to set the following threshold conditions for inodes:

Nearly Full

Specifies the percentage at which a volume is considered to have consumed most of its inodes:

- Default value: 80 percent
- Event generated: Inodes Nearly Full
- Event severity: Warning

Full

Specifies the percentage at which a volume is considered to have consumed all of its inodes:

- Default value: 90 percent
- Event generated: Inodes Full
- Event severity: Error

Command buttons

The command buttons enable you to save or cancel the setup options:

Restore to Factory Defaults

Enables you to restore the configuration settings to the factory default values.

Save

Saves the configuration settings for the selected option.

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Configuring global volume threshold values on page 127

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 user quota events are generated. The OnCommand Administrator can configure alerts that notify the specified recipients of the user or user group quota events.

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

Viewing user and user group quotas

The Storage Virtual Machine 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, soft and hard limits set on the disks and files, used disk and file space, and disk threshold value.

Before you begin

You must have one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. Select the SVM for which you want to view the user and user group quota details.
3. Click **User and Group Quotas**.

Related concepts

Overview of the quota process on page 139

Related tasks

Adding a user on page 379

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.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have reviewed the *guidelines for creating rules* on page 141.

About this task

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 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

Steps

1. Click **Health** > **Administration** > **Setup Options**.
2. In the **Setup Options** dialog box, click **Quota Settings** > **Email Address Rules**.
3. In the **Rules to Generate User and Group Quota Email Address** dialog box, enter the rule based on your criteria.
4. 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.
5. Click **Save and Close**.
6. Verify that the email address you created is displayed in the **User and Group Quotas** tab of the **Storage Virtual Machine details** page.

Related tasks

Editing user and group quota email addresses on page 138
Creating an email notification format for user and user group quotas on page 138

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).

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Quota Settings** > **Email Notification Format**.
4. Enter or modify the details in the **From**, **Subject**, and **Email Details** fields.
5. Click **Preview** to preview the email notification.

   The Notification Preview dialog box is displayed with the email notification.
6. Click **Close**.
7. Modify the content of the email notification, if required.
8. Click **Save**.

Related tasks

Creating rules to generate email addresses on page 137

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.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have reviewed the *guidelines for creating rules* on page 141.

About this task

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. Click **Storage** > **Storage Virtual Machines**.

2. Select the SVM for which you want to view the user and user group quota details.
3. Click **User and Group Quotas**, and then click **Edit Email Address**.
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	**a.** Delete the email address in the **Email Address** field. **b.** Click **Save**. **c.** Refresh the browser by pressing F5 to reload the Storage Virtual Machine details page. 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	**a.** Modify the email address in the **Email Address** field. **b.** Click **Save**. The rules to generate the user and group quota email addresses are no longer applicable to the quota.

Related tasks

Creating rules to generate email addresses on page 137

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 Data ONTAP to send a notification when specified thresholds are exceeded, and hard quotas prevent a write operation from succeeding when specified thresholds are exceeded.

When Data 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 will fail 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 will succeed 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.

Related concepts

About quotas on page 140
Why you use quotas on page 140

About quotas

Quotas provide a way to restrict or track the disk space and number of files used by a user, group, or qtree.

Why you use quotas

You can use quotas to limit resource usage, 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 Setup Options dialog box 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 dialog box

The Email Notification Format dialog box displays the notification format 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 can be modified. By default, this is the email address that is specified under Notification in the General Settings area of the Setup Options dialog box.

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.

Save and Close

Saves the changes made to the notification format and closes the Email Notification Format dialog box.

Cancel

Cancels the recent changes and closes the Email Notification Format dialog box.

Related references

Email Notification Format dialog box on page 140
Setup Options dialog box on page 41
Rules to Generate User and Group Quota Email Address dialog box on page 141

Rules to Generate User and Group Quota Email Address dialog box

The Rules to Generate User and Group Quota Email Address dialog box enables you to create rules to specify email addresses 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 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 variable operator* *) 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 variable operator* '*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.

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.

Save and Close

Validates the syntax of the rule, saves the rule if there are no errors, and closes the Rules to Generate User and Group Quota Email Address dialog box.

Cancel

Cancels the recent changes and closes the Rules to Generate User and Group Quota Email Address dialog box.

Related tasks

Creating rules to generate email addresses on page 137

Related references

Setup Options dialog box on page 41
Email Notification Format dialog box on page 140

Managing and monitoring clusters and cluster objects

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

Understanding cluster monitoring

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

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

The following flowchart illustrates the Unified Manager monitoring process:

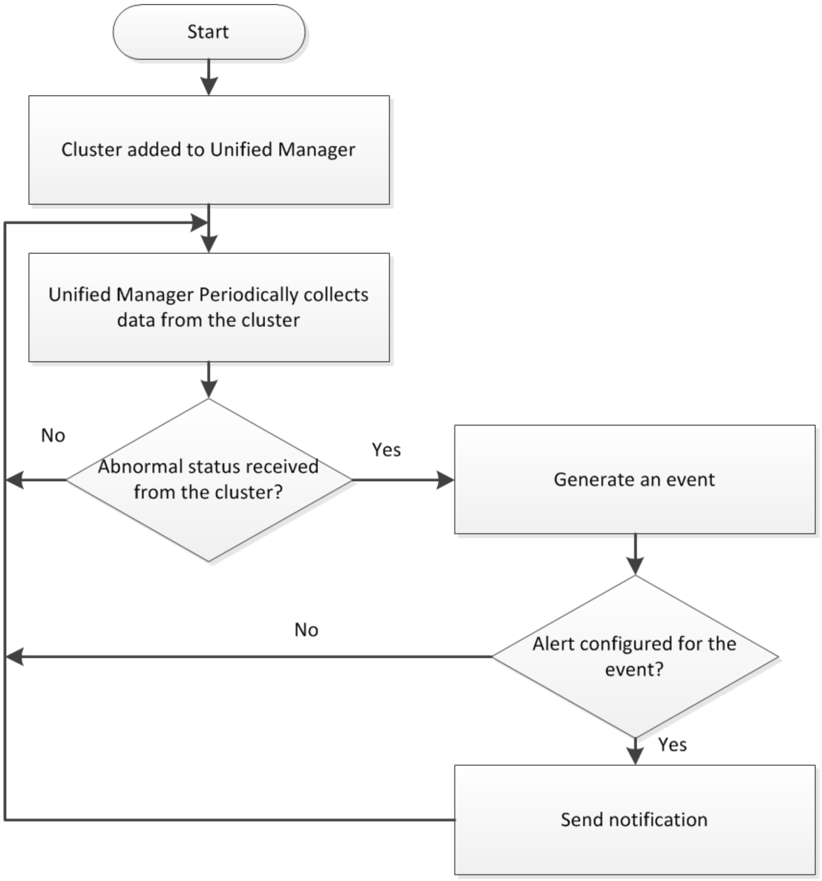
Start
Cluster added to Unified Manager
Unified Manager Periodically collects data from the cluster
No
Abnormal status received from the cluster?
Yes
Generate an event
No
Alert configured for the event?
Yes
Send notification

Understanding node root volumes

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

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

Understanding capacity events and thresholds for node root aggregates

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

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

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

Understanding quorum and epsilon

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

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

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

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

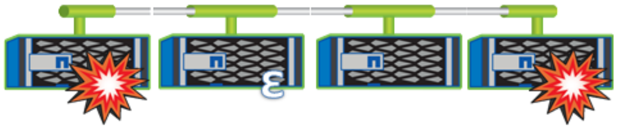

Epsilon is automatically assigned to the first node when the cluster is created. If the node that holds epsilon becomes unhealthy, takes over its high-availability partner, or is taken over by its high-availability partner, then epsilon is automatically reassigned to a healthy node in a different HA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

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

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

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

Removing a cluster in case of discovery failure

If the discovery of a cluster fails, you can remove the cluster by using the Manage Data Sources page.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The **Remove** button in the Manage Data Sources page is enabled only when the discovery of the selected cluster fails.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Data Sources**.
3. Select the cluster that failed to be discovered.

   You can use the filter on the State column to list the clusters for which discovery failed.
4. Click **Remove**.

Related tasks

Adding clusters on page 20

Related references

Manage Data Sources page on page 229

Restarting cluster discovery

If a cluster discovery process fails, you can restart the process by using the Manage Data Sources page. For example, you might want to restart cluster discovery if there were network issues while you were adding a cluster to the Unified Manager database.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Data Sources**.
3. Select the name of the cluster that failed to be discovered.

   You can use the filter on the State column to list the clusters for which discovery failed.

4. Click **Rediscover**.

Related references

Manage Data Sources page on page 229

Viewing the cluster list and details

You can use the Clusters page to view your inventory of clusters. You can also view details such as the cluster health, capacity, configuration, LIFs, and nodes in that cluster by using the Cluster details page.

Before you begin

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

About this task

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

Steps

1. Click **Storage** > **Clusters**.
2. View the complete details of the cluster by clicking the cluster name.

Related tasks

Adding a user on page 379
Exporting data to CSV files on page 31
Adding an alert on page 102
Editing clusters on page 22

Related references

Clusters page on page 168
Cluster details page on page 170

Viewing the node list and details

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

Before you begin

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

About this task

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

Steps

1. Click **Storage** > **Clusters**.
2. On the **Clusters** page, click the cluster name whose node details you want to view.
3. On the **Cluster details** page, click the **Nodes** tab.

   The left pane displays the list of HA pairs. By default, the HA Details tab is open, which displays HA state details and events related to the selected HA pair.
4. To view other details about the node, perform the appropriate action:

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

Related tasks

Adding a user on page 379
Viewing the cluster list and details on page 147
Exporting data to CSV files on page 31
Adding an alert on page 102
Editing clusters on page 22

Related references

Nodes page on page 182

Viewing the SVM list and details

From the Storage Virtual Machines page, you can monitor your inventory of Storage Virtual Machines (SVMs). You can use the Storage Virtual Machine details page to view detailed information about SVMs that are monitored.

Before you begin

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

About this task

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

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. Choose one of the following ways to view the SVM details:
   - To view minimal details, position the cursor over the SVM name.
   - To view the complete details, click the SVM name.
     You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. Optional: View the objects related to the SVM by clicking **View Related** in the minimal details dialog box.

Related tasks

Adding a user on page 379
Exporting data to CSV files on page 31
Adding an alert on page 102
Viewing the details of SVMs with Infinite Volume on page 360

Related references

Storage Virtual Machines page on page 184
Storage Virtual Machine details page on page 185

Viewing the aggregate list and details

From the Aggregates page, you can monitor your inventory of aggregates. You can use the Aggregate details page to view detailed information about aggregates that are monitored.

Before you begin

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

About this task

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

Steps

1. Click **Storage** > **Aggregates**.
2. Choose one of the following ways to view the aggregate details:
    - To view minimal details, position the cursor over the aggregate name.
    - To view the complete details, click the aggregate name.
      You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. Optional: View the objects related to the aggregate by clicking **View Related** from the minimal details dialog box.

Related tasks

Adding a user on page 379
Exporting data to CSV files on page 31
Adding an alert on page 102
Editing aggregate threshold settings on page 128

Related references

Aggregates page on page 200
Aggregate details page on page 203

Viewing storage pool details

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

Before you begin

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

Steps

1. Click **Storage** > **Aggregates**.
2. Click the required aggregate name.

   The details of the selected aggregate are displayed.
3. Click the **Disks** tab.
4. In the Cache table, move the pointer over the name of the required storage pool.

   The details of the storage pool are displayed.

Viewing the volume list and details

From the Volumes page, you can monitor your inventory of volumes. You can use the Volume details page to view detailed information about volumes that are monitored, including the capacity, efficiency, and configuration of the volumes. You can also view information about the related devices and related alerts for a specific volume.

Before you begin

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

Steps

1. Click **Storage** > **Volumes**.
2. Choose one of the following ways to view the volume details:
   - To view minimal details, position the cursor over the volume name.
   - To view the complete details, click the volume name.
     You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. Optional: View the objects related to the volume by clicking **View Related** from the minimal details dialog box.

Related tasks

Adding a user on page 379
Exporting data to CSV files on page 31
Adding an alert on page 102
Editing volume threshold settings on page 129

Related references

Volumes page on page 211
Volume details page on page 216

Viewing the CIFS shares

You can use the Storage Virtual Machine details page to view detailed information about the CIFS share hosted by the Storage Virtual Machine (SVM). You can view details such as the share name, junction path, containing objects, security settings, and export policies defined for the share.

Before you begin

- CIFS license must be enabled on the cluster.
- LIFs serving the CIFS shares must be configured.
- You must have the Operator, OnCommand Administrator, or Storage Administrator role.

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. Select the SVM for which you want to view the CIFS share details.
3. Click **CIFS Shares**.

Related concepts

About CIFS and SMB on page 163

Related tasks

Troubleshooting access to CIFS shares on page 464

Viewing the data sources

You can use the Manage Data Sources page to view detailed information about the data sources and clusters that are added to the Unified Manager database. You can view details such as the discovery status, data source name and type, supported discovery operations, operation state, operation start and end time, and description of the operation.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Data Sources**.

   Discovery details about the data sources are displayed in the Manage Data Sources page.

Viewing the list of Snapshot copies

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

Before you begin

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

Steps

1. Click **Storage** > **Volumes**.
2. In the **Volumes** page, select the appropriate volume that contains the Snapshot copies you want to view.
3. In the **Capacity** tab of the **Volume details** page, click the link next to **Snapshot Copies**.

   For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. For volumes in a cluster running Data ONTAP 8.1.x, you should click the **View** link.

   The list of Snapshot copies is displayed.

Related references

Volumes page on page 211
Volume details page on page 216

Deleting Snapshot copies

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

The volume must be online.

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

About this task

- You cannot delete the base Snapshot copy in a parent volume if a FlexClone volume is using that Snapshot copy.
  The base Snapshot copy is the Snapshot copy that is used to create the FlexClone volume and displays the status `Busy` and Application Dependency as `Busy,Vclone` in the parent volume.
- You cannot delete a locked Snapshot copy that is used in a SnapMirror relationship.
  The Snapshot copy is locked and is required for the next update.

Steps

1. Click **Storage** > **Volumes**.
2. In the **Volumes** page, select the volume that contains the Snapshot copies you want to delete.
3. In the **Capacity** tab of the **Volume details** page, click the link next to **Snapshot Copies**.

   For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. For volumes in a cluster running Data ONTAP 8.1.x, you should click the **View** link.

   The list of Snapshot copies on the volume is displayed.
4. Select the Snapshot copies you want to delete.
5. Click **Delete Selected**.

Calculating reclaimable space for Snapshot copies

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

Before you begin

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

The volume must be online.

Steps

1. Click **Storage** > **Volumes**.
2. In the **Volumes** page, select the appropriate volume for which you want calculate the reclaimable space.
3. In the **Capacity** tab of the **Volume details** page, click the link next to **Snapshot Copies**.

   For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. For volumes in a cluster running Data ONTAP 8.1.x, you should click the **View** link.

   The list of Snapshot copies on the volume is displayed.
4. Select the Snapshot copies for which you want to calculate the reclaimable space.
5. Click **Calculate**.

   The reclaimable space (in percentage, and KB, MB, GB, and so on) on the volume is displayed.
6. To recalculate the reclaimable space, select the required Snapshot copies and click **Recalculate**.

Understanding clusters and cluster objects

You should understand some of the basic concepts of clusters and cluster objects to monitor and manage your clusters.

Related concepts

Types of workloads monitored by Performance Manager on page 60

Related references

Workload performance measurement values on page 61

What a cluster is

A cluster consists of one or more nodes grouped together as (HA pairs) to form a scalable cluster. Creating a cluster enables the nodes to pool their resources and distribute work across the cluster, while presenting administrators with a single entity to manage. Clustering also enables continuous service to end users if individual nodes go offline.

- The maximum number of nodes within a cluster depends on the platform model and licensed protocols.
- Each node in the cluster can view and manage the same volumes as any other node in the cluster.
  The total file-system namespace, which comprises all of the volumes and their resultant paths, spans the cluster.
- The nodes in a cluster communicate over a dedicated, physically isolated and secure Ethernet network.
  The cluster logical interfaces (LIFs) on each node in the cluster must be on the same subnet.
- When new nodes are added to a cluster, there is no need to update clients to point to the new nodes.
  The existence of the new nodes is transparent to the clients.
- If you have a two-node cluster (a single HA pair), you must configure cluster high availability (HA).
- You can create a cluster on a stand-alone node, called a single-node cluster.
  This configuration does not require a cluster network, and enables you to use the cluster ports to serve data traffic. However, nondisruptive operations are not supported on single-node clusters.

Related concepts

Understanding clusters and cluster objects on page 153

What SVMs are

Storage Virtual Machines (SVMs, formerly known as Vservers) contain data volumes and one or more LIFs through which they serve data to the clients. Starting with clustered Data ONTAP 8.1.1, SVMs can either contain one or more FlexVol volumes, or a single Infinite Volume.

SVMs securely isolate the shared virtualized data storage and network, and each SVM appears as a single dedicated server to the clients. Each SVM has a separate administrator authentication domain and can be managed independently by its SVM administrator.

In a cluster, SVMs facilitate data access. A cluster must have at least one SVM to serve data. SVMs use the storage and network resources of the cluster. However, the volumes and LIFs are exclusive to the SVM. Multiple SVMs can coexist in a single cluster without being bound to any node in a cluster. However, they are bound to the physical cluster on which they exist.

A cluster can have one or more SVMs with FlexVol volumes and SVMs with Infinite Volume.

SVM with FlexVol volumes

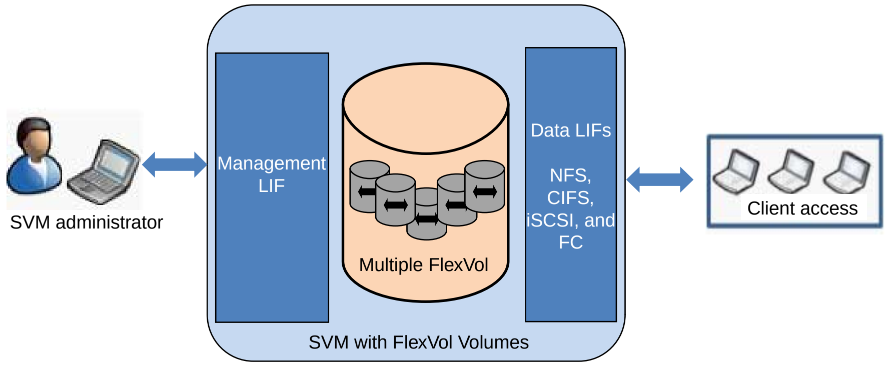

Each SVM with FlexVol volumes in a NAS environment presents a single directory hierarchical view and has a unique namespace. The namespace enables NAS clients to access data without specifying the physical location of the data. The namespace also enables the cluster and SVM administrators to manage distributed data storage as a single directory with multiple levels of hierarchy.

The volumes within each NAS SVM are related to each other through junctions and are mounted on junction paths. These junctions present the file system in each volume. The root volume of the SVM is a FlexVol volume that resides at the top level of the namespace hierarchy; additional volumes are mounted to the SVM root volume to extend the namespace. As volumes are created for the SVM, the root volume of the SVM contains junction paths.

SVMs with FlexVol volumes can contain files and LUNs. They provide file-level data access by using NFS and CIFS protocols for the NAS clients, and block-level data access by using iSCSI and Fibre Channel (FC) (FCoE included) for SAN hosts.

SVM with Infinite Volume

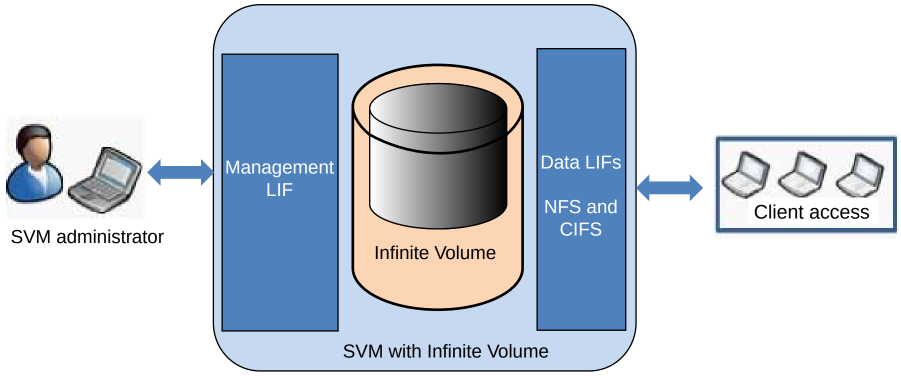

SVMs with Infinite Volume can contain only one Infinite Volume to serve data. Each SVM with Infinite Volume includes only one junction path, which has a default value of `/NS`. The junction provides a single mount point for the large namespace provided by the SVM with Infinite Volume. You cannot add more junctions to an SVM with Infinite Volume. However, you can increase the size of the Infinite Volume.

SVMs with Infinite Volume can contain only files. They provide file-level data access by using NFS and CIFS protocols. SVMs with Infinite Volume cannot contain LUNs and do not provide block-level data access.

Note: The Data ONTAP command-line interface (CLI) continues to use the term *Vserver* in the output, and `vserver` as a command or parameter name has not changed.

Why you use SVMs

Storage Virtual Machines (SVMs, formerly known as Vservers) provide data access to clients regardless of the physical storage or controller, similar to any storage system. SVMs provide benefits such as nondisruptive operations, scalability, security, and unified storage.

SVMs provide the following benefits:

- Multi-tenancy

  SVM is the fundamental unit of secure multi-tenancy, which enables partitioning of the storage infrastructure so that it appears as multiple independent storage systems. These partitions isolate the data and management.
- Nondisruptive operations

  SVMs can operate continuously and nondisruptively for as long as they are needed. SVMs help clusters to operate continuously during software and hardware upgrades, addition and removal of nodes, and all administrative operations.
- Scalability

  SVMs meet on-demand data throughput and the other storage requirements.
- Security

  Each SVM appears as a single independent server, which enables multiple SVMs to coexist in a cluster while ensuring no data flows among them.
- Unified storage

  SVMs can serve data concurrently through multiple data access protocols. SVMs provide file-level data access through NAS protocols, such as CIFS and NFS, and block-level data access through SAN protocols, such as iSCSI and FC (FCoE included). SVMs can serve data to SAN and NAS clients independently at the same time.

  Note: SVMs with Infinite Volume can serve data only through NFS and CIFS protocols.
- Delegation of management

  Each SVM can have its own user and administration authentication. SVM administrators can manage the SVMs that they are authorized to access. However, SVM administrators have privileges assigned by the cluster administrators.
- Easy management of large datasets

  With SVMs with Infinite Volume, management of large and unstructured data is easier because the SVM administrator can manage one data container instead of many.

What logical storage is

Logical storage refers to the storage resources provided by Data ONTAP that are not tied to a physical resource.

Logical storage resources are associated with a Storage Virtual Machine (SVM, formerly known as Vserver), and they exist independently of any specific physical storage resource such as a disk, array LUN, or aggregate. Logical storage resources include volumes of all types and qtrees, as well as the capabilities and configurations you can use with these resources, such as Snapshot copies, deduplication, compression, and quotas.

For more information about SVMs, see the *System Administration Reference*.

How volumes work

Volumes are data containers that enable you to partition and manage your data. Understanding the types of volumes and their associated capabilities enables you to design your storage architecture for maximum storage efficiency and ease of administration.

Volumes are the highest-level logical storage object. Unlike aggregates, which are composed of physical storage resources, volumes are completely logical objects.

Data ONTAP provides two types of volumes: FlexVol volumes and Infinite Volumes. There are also volume variations, such as FlexClone volumes, FlexCache volumes, SnapLock volumes, data protection mirrors, and load-sharing mirrors. Not all volume variations are supported for both types of volumes. Data ONTAP efficiency capabilities, compression and deduplication, are supported for both types of volumes.

Volumes contain file systems in a NAS environment, and LUNs in a SAN environment.

Volumes are associated with one Storage Virtual Machine (SVM). The SVM is a virtual management entity, or server, that consolidates various cluster resources into a single manageable unit. When you create a volume, you specify the SVM it is associated with. The type of the volume (FlexVol volume or Infinite Volume) is determined by an immutable SVM attribute.

Volumes have a language. The language of the volume determines the character set Data ONTAP uses to display file names and data for that volume. The default value for the language of the volume is the language of the SVM.

Volumes depend on their associated aggregates for their physical storage; they are not directly associated with any concrete storage objects, such as disks or RAID groups. If the cluster administrator has assigned specific aggregates to an SVM, then only those aggregates can be used to provide storage to the volumes associated with that SVM. This impacts volume creation, and also copying and moving FlexVol volumes between aggregates.

For more information about Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

For more information about SVMs, see the *System Administration Reference*.

For more information about data protection mirrors, see the *Clustered Data ONTAP Data Protection Guide*.

For more information about physical storage resources such as aggregates, disks, and RAID groups, see the *Clustered Data ONTAP Physical Storage Management Guide*.

Considerations for using thin-provisioned FlexVol volumes

You can configure your thin-provisioned volumes so that they appear to provide more storage than they have available, provided that the storage that is actually being used does not exceed the available storage. However, you should understand how thin-provisioned volumes can act differently from fully provisioned volumes.

If the volumes associated with an aggregate show more storage as available than the physical resources available to that aggregate, the aggregate is *overcommitted*. When an aggregate is overcommitted, it is possible for writes to LUNs or files in volumes contained by that aggregate to fail if there is not sufficient free space available to accommodate the write.

If you have overcommitted your aggregate, you must monitor your available space and add storage to the aggregate as needed to avoid write errors due to insufficient space.

Aggregates can provide storage to FlexVol volumes associated with more than one Storage Virtual Machine (SVM). When sharing aggregates for thin-provisioned FlexVol volumes in a multi-tenancy environment, be aware that one tenant's aggregate space availability can be adversely affected by the growth of another tenant's volumes.

Volume states

Volumes can be in one of four states—online, offline, restricted, or mixed.

On the Volume Details page, the volume status is displayed in parentheses at the top of the page next to the volume name.

The following table displays the possible states for volumes.

State	Description
online	Read and write access to this volume is allowed.
offline	No access to the volume is allowed.
restricted	Some operations, such as parity reconstruction, are allowed, but data access is not allowed.
mixed	The constituents of an Infinite Volume are not all in the same state.

What LIFs are

A LIF (logical interface) is an IP address or WWPN with associated characteristics, such as a role, a home port, a home node, a list of ports to fail over to, and a firewall policy. You can configure LIFs on ports over which the cluster sends and receives communications over the network.

LIFs can be hosted on the following ports:

- Physical ports that are not part of interface groups
- Interface groups
- VLANs
- Physical ports or interface groups that host VLANs

While configuring SAN protocols such as FC on a LIF, it will be associated with a WWPN.

For more information about configuring WWPNs for LIFs using the FC protocol, see the *Clustered Data ONTAP SAN Administration Guide*.

The following figure illustrates the port hierarchy in a clustered Data ONTAP system:

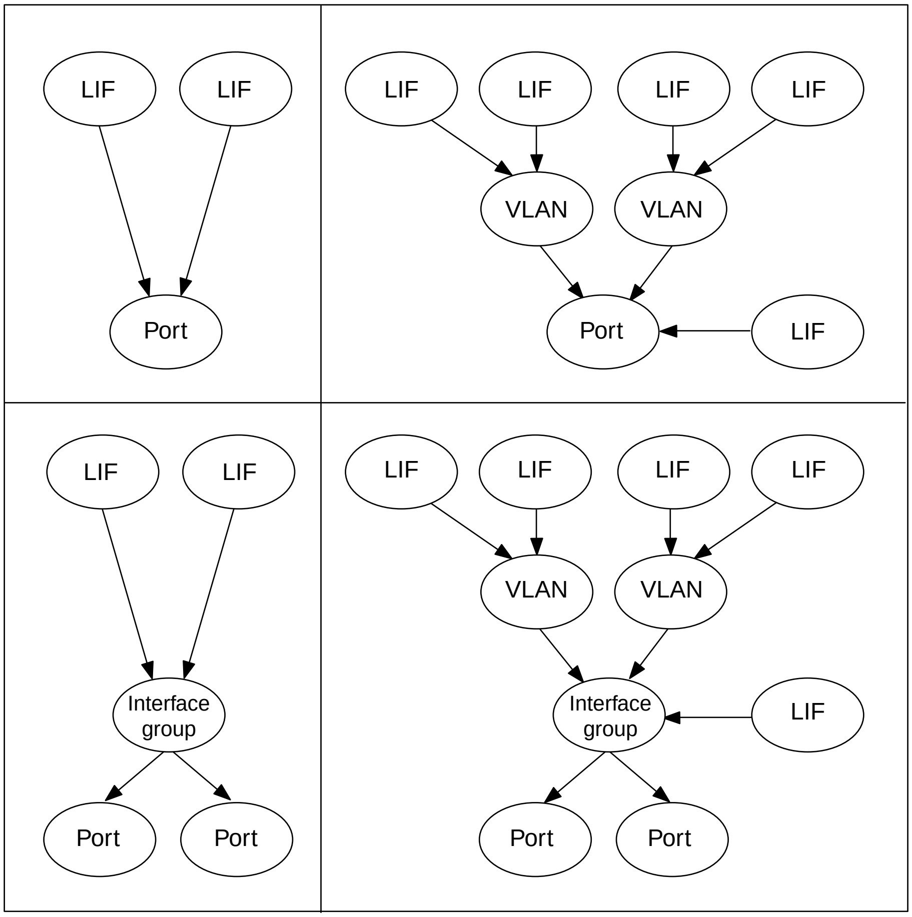

Types of network ports

Ports are either physical ports (NICs) or virtualized ports such as interface groups or VLANs. Interface groups treat several physical ports as a single port, while VLANs subdivide a physical port into multiple separate logical ports.

physical ports

LIFs can be configured directly on physical ports.

interface group

A port aggregate containing two or more physical ports that act as a single trunk port. An interface group can be single-mode, multimode, or dynamic multimode.

VLAN

A logical port that receives and sends VLAN-tagged (IEEE 802.1Q standard) traffic. VLAN port characteristics include the VLAN ID for the port. The underlying physical port or interface group ports are considered VLAN trunk ports, and the connected switch ports must be configured to trunk the VLAN IDs.

The underlying physical port or interface group ports for a VLAN port can continue to host LIFs, which transmit and receive untagged traffic.

The port naming convention is e<number>letter:

- The first character describes the port type.
“e” represents Ethernet.
- The second character indicates the slot in which the port adapter is located.

- The third character indicates the port's position on a multiport adapter.
  "a" indicates the first port, "b" indicates the second port, and so on.

For example, "e0b" indicates that an Ethernet port is the second port on the node's motherboard.

VLANs must be named by using the syntax port_name-vlan-id. "port_name" specifies the physical port or interface group and "vlan-id" specifies the VLAN identification on the network. For example, "e1c-80" is a valid VLAN name.

What the Snapshot copy reserve is

The Snapshot copy reserve sets a specific percent of the disk space for Snapshot copies. By default, the Snapshot copy reserve is 5 percent of the disk space for a FlexVol volume and 0 percent for aggregates.

The active file system cannot consume the Snapshot copy reserve space, but the Snapshot copy reserve, if exhausted, can use space in the active file system.

What an HA pair is

An HA pair is two storage systems (nodes) whose controllers are connected to each other directly. In this configuration, one node can take over its partner's storage to provide continued data service if the partner goes down.

You can configure the HA pair so that each node in the pair shares access to a common set of storage, subnets, and tape drives, or each node can own its own distinct set of storage.

The controllers are connected to each other through an interconnect. This allows one node to serve data that resides on the disks of its failed partner node. Each node continually monitors its partner, mirroring the data for each other's nonvolatile memory (NVRAM or NVMEM). The interconnect is internal and requires no external cabling if both controllers are in the same chassis.

Takeover is the process in which a node takes over the storage of its partner. *Giveback* is the process in which that storage is returned to the partner. Both processes can be initiated manually or configured for automatic initiation.

How HA pairs relate to the cluster

HA pairs are components of the cluster, and both nodes in the HA pair are connected to other nodes in the cluster through the data and cluster networks. But only the nodes in the HA pair can take over each other's storage.

Although the controllers in an HA pair are connected to other controllers in the cluster through the cluster network, the HA interconnect and disk-shelf connections are found only between the node and its partner and their disk shelves or array LUNs.

The HA interconnect and each node's connections to the partner's storage provide physical support for high-availability functionality. The high-availability storage failover capability does not extend to other nodes in the cluster.

> **Note:** Network failover does not rely on the HA interconnect and allows data network interfaces to failover to different nodes in the cluster outside the HA pair. Network failover is different than storage failover since it enables network resiliency across all nodes in the cluster.

Non-HA (or stand-alone) nodes are not supported in a cluster containing two or more nodes. Although single-node clusters are supported, joining two separate single-node clusters to create one cluster is not supported, unless you wipe clean one of the single-node clusters and join it to the other to create a two-node cluster that consists of an HA pair.

The following diagram shows two HA pairs. The multipath HA storage connections between the nodes and their storage are shown for each HA pair. For simplicity, only the primary connections to the data and cluster networks are shown.

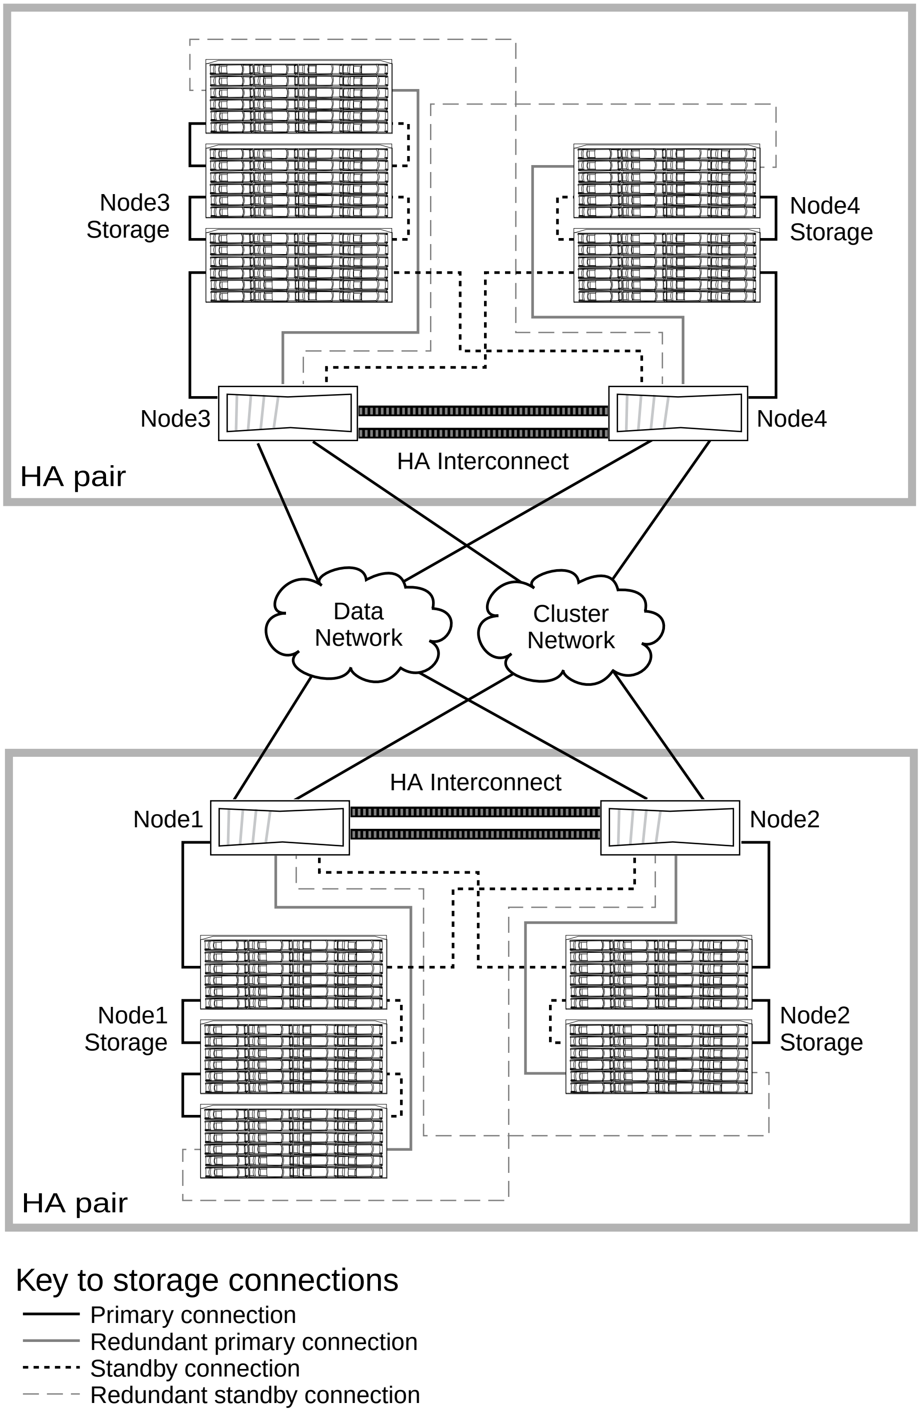

Possible storage failover scenarios in this cluster are as follows:

- Node1 fails and Node2 takes over Node1's storage.
- Node2 fails and Node1 takes over Node2's storage.
- Node3 fails and Node4 takes over Node3's storage.
- Node4 fails and Node3 takes over Node4's storage.

If Node1 *and* Node2 both fail, the storage owned by Node1 and Node2 becomes unavailable to the data network. Although Node3 and Node4 are clustered with Node1 and Node2, they do not have direct connections to Node1 and Node2's storage and cannot take over their storage.

When takeovers occur

You can initiate takeovers manually or they can occur automatically when a failover event happens, depending on how you configure the HA pair. In some cases, takeovers occur automatically, regardless of configuration.

Takeovers can occur under the following conditions:

- When you manually initiate takeover
- When a node in an HA pair with the default configuration for immediate takeover on panic undergoes a software or system failure that leads to a panic

  By default, the node automatically performs a giveback, returning the partner to normal operation after the partner has recovered from the panic and booted up.
- When a node in an HA pair undergoes a system failure (for example, a loss of power) and cannot reboot

  Note: If the storage for a node also loses power at the same time, a standard takeover is not possible.
- When a node does not receive heartbeat messages from its partner

  This could happen if the partner experienced a hardware or software failure that did not result in a panic but still prevented it from functioning correctly.
- When hardware-assisted takeover is enabled and it triggers a takeover when the remote management device (Service Processor) detects failure of the partner node

Connections and components of an HA pair

Each node in an HA pair requires a network connection, an HA interconnect between the controllers, and connections to both its own disk shelves and its partner node's shelves.

The following diagram shows a standard HA pair with native DS4243 disk shelves and multipath HA:

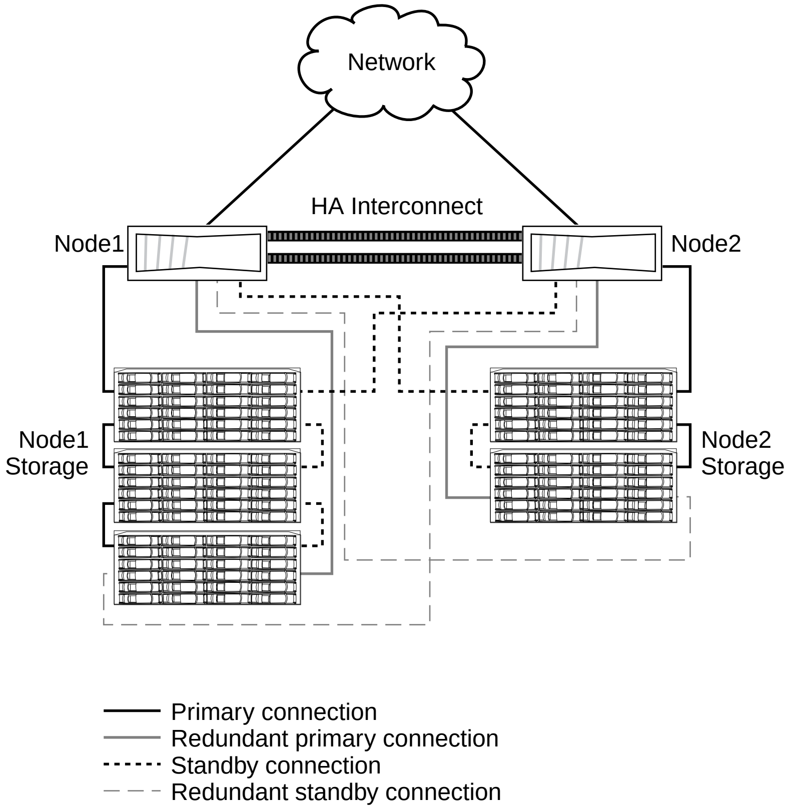

Understanding takeover and giveback

Takeover and giveback are the operations that let you take advantage of the HA configuration to perform nondisruptive operations and avoid service interruptions. Takeover is the process in which a node takes over the storage of its partner. Giveback is the process in which the storage is returned to the partner.

About CIFS and SMB

Data ONTAP supports all of the most common file protocols, including the CIFS protocol to enable file sharing from host storage systems. When your system is first installed and CIFS is configured in Workgroup mode, a login named "administrator" is automatically created. You can use this login to access shares with a blank password.

The CIFS protocol is used to share files. CIFS is the method of transport for Windows Shares. CIFS is an extension of the Server Message Block (SMB) protocol, which is a file-sharing protocol used on Windows and UNIX systems. SMB runs over several different types of networks, including TCP/IP. For most purposes, SMB is superseded by CIFS.

Related concepts

Supported SMB clients and domain controllers on page 163

Supported SMB clients and domain controllers

Before you can use SMB with your Storage Virtual Machine (SVM), you need to know which SMB clients and domain controllers Data ONTAP supports.

For the latest information about which SMB clients and domain controllers Data ONTAP supports, see the Interoperability Matrix.

mysupport.netapp.com/matrix

What a Snapshot copy is

A Snapshot copy is a frozen, read-only image of a flexible volume, or an aggregate that captures the state of the file system at a point in time. Snapshot copies are your first line of defense to backup and restore data.

When Snapshot copies are created, Data ONTAP maintains a configurable Snapshot copy schedule that creates and deletes Snapshot copies automatically for each volume. You can also create and delete Snapshot copies manually.

You can store up to 255 Snapshot copies at one time on each volume.

You can specify the percentage of disk space that Snapshot copies can occupy. The default space reserved for Snapshot copies is zero percent for SAN and VMware volumes. For NAS volumes, it is five percent on storage systems running Data ONTAP 8.1.

Storage QoS

Data ONTAP 8.2 introduces Storage QoS, which can help you manage the risks that accompany meeting performance objectives for workloads. You can use Storage QoS to limit the throughput to workloads to a Storage Virtual Machine (SVM), or to groups of volumes or LUNs within an SVM, and to monitor IOPS and MBps performance.

You can reactively limit workload performance to ensure fair resource usage. If you have a cloud infrastructure, you might proactively limit workloads, as defined by their service levels.

For example, you can prevent runaway workloads from impacting other workloads in a shared storage infrastructure by applying a throughput limit to the runaway workload. In a service-provider environment, you can proactively set throughput limits at an SVM level, where an SVM maps to a tenant. This throughput limit ensures a consistent performance for each tenant as you add more tenants to the shared storage infrastructure and prevents tenants from affecting each other's workload performance.

Storage QoS is supported on clusters that have up to eight nodes.

For information about how to use Storage QoS, see the *System Administration Reference*.

How the maximum throughput limit works

You can specify one service-level objective for a Storage QoS policy group: a maximum throughput limit. A maximum throughput limit, which you define in terms of IOPS or MBps, specifies the throughput that the workloads in the policy group cannot collectively exceed.

When you specify a maximum throughput for a policy group, Storage QoS controls client operations to ensure that the combined throughput for all workloads in the policy group does not exceed the specified maximum throughput.

For example, assume that you create the policy group “untested_apps” and specify a maximum throughput of 300 MBps. You assign three volumes to the policy group. The combined throughput to those three volumes cannot exceed 300 MBps.

Note: The combined throughput to the workloads in a policy group might exceed the specified limit by up to 10%. A deviation might occur if you have a workload that experiences rapid changes in throughput (sometimes called a *bursty workload*).

Note the following about specifying a maximum throughput:

- A throughput limit applies to all clients that access a storage object.
- Do not set the limit too low, because you might underutilize the cluster.
- Consider the minimum amount of throughput that you want to reserve for workloads that do not have limits.

For example, you can ensure that your critical workloads get the throughput that they need by limiting noncritical workloads.

- You might want to provide room for growth.
  For example, if you see an average utilization of 500 IOPS, you might specify a limit of 1,000 IOPS.

How throttling a workload can affect non-throttled workload requests from the same client

In some situations, throttling a workload (I/O to a storage object) can affect the performance of non-throttled workloads if the I/O requests are sent from the same client.

If a client sends I/O requests to multiple storage objects and some of those storage objects belong to Storage QoS policy groups, performance to the storage objects that do not belong to policy groups might be degraded. Performance is affected because resources on the client, such as buffers and outstanding requests, are shared.

For example, this might affect a configuration that has multiple applications or virtual machines running on the same host.

This behavior is likely to occur if you set a low maximum throughput limit and there are a high number of I/O requests from the client.

If this occurs, you can increase the maximum throughput limit or separate the applications so they do not contend for client resources.

Controlling and monitoring I/O performance to FlexVol volumes by using Storage QoS

You can control input/output (I/O) performance to FlexVol volumes by assigning volumes to Storage QoS policy groups. You might control I/O performance to ensure that workloads achieve specific performance objectives or to throttle a workload that negatively impacts other workloads.

About this task

Policy groups enforce a maximum throughput limit (for example, 100 MB/s). You can create a policy group without specifying a maximum throughput, which enables you to monitor performance before you control the workload.

You can also assign Storage Virtual Machines (SVMs) with FlexVol volumes, LUNs, and files to policy groups.

Note the following requirements about assigning a volume to a policy group:

- The volume must be contained by the SVM to which the policy group belongs.
  You specify the SVM when you create the policy group.
- If you assign a volume to a policy group, then you cannot assign the volume's containing SVM or any child LUNs or files to a policy group.

For more information about how to use Storage QoS, see the *System Administration Reference*.

Steps

1. Use the `qos policy-group create` command to create a policy group.
2. Use the `volume create` command or the `volume modify` command with the `-qos-policy-group` parameter to assign a volume to a policy group.
3. Use the `qos statistics` commands to view performance data.
4. If necessary, use the `qos policy-group modify` command to adjust the policy group's maximum throughput limit.

What network processing is

Network processing is a software component in the cluster that handles read and write requests between client applications and the network protocols on the cluster.

The network processing operations communicate with the data processing software component to locate the storage aggregate that will fulfill the request. Once the request is fulfilled, the data processing component communicates with the network processing component to transfer the information to the requesting client.

What data processing is

Data processing is a software component in the cluster that handles read and write requests to the target storage aggregate.

The data processing component receives read and write requests from the network processing software component. It locates the storage aggregate that can fulfill the requests, and communicates with the network processing component to transfer information to the requesting client applications.

Storage aggregates and disks

A storage aggregate is a logical grouping of physical storage that is protected by RAID technology. When you create a storage aggregate, you specify the number of disks in the underlying RAID set and the level of RAID protection for the set (RAID-4 or RAID-DP).

A storage aggregate contains a defined amount of physical storage that can be expanded dynamically at any time. For example, if a storage aggregate is made up of three disks, an additional disk can be easily added to the storage aggregate, and the logical size of the storage aggregate can be increased accordingly, without disrupting any user or process currently using that storage aggregate. Storage aggregates can include hot spare disks that hold data if one of the other disks fails.

Aggregate states

The state of an aggregate indicates its availability or whether it is involved in a specific process.

State	Description
Offline	Read or write access is not allowed.
Restricted	Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
Online	Read and write access to volumes hosted on this aggregate is allowed.
Creating	The aggregate is being created.
Destroying	The aggregate is being destroyed.
Failed	The aggregate cannot be brought online.
Frozen	The aggregate is (temporarily) not serving requests.
Inconsistent	The aggregate has been marked corrupted; contact technical support.
Iron Restricted	Diagnostic tools cannot be run on the aggregate.
Mounting	The aggregate is being mounted.

State	Description
Partial	At least one disk was found for the aggregate, but two or more disks are missing.
Quiescing	The aggregate is being quiesced.
Quiesced	The aggregate is quiesced.
Reverted	The revert of an aggregate is completed.
Unmounted	The aggregate is offline.
Unmounting	The aggregate is being taken offline.
Unknown	The aggregate is discovered, but the aggregate information is not yet retrieved by the OnCommand application server.

Aggregate capacity states

The storage capacity of an aggregate can be in 1 of 3 states: normal, warning, or error. The states are based on pre-defined capacity thresholds.

State	Description
✓ Normal	Used capacity is under the Warning and Error thresholds.
⚠ Warning	Used capacity is above the Warning threshold of 85% of the total aggregate capacity.
! Error	Used capacity is above the Error threshold of 95% of the total aggregate capacity.

Using thin provisioning with FlexVol volumes

Using thin provisioning, you can appear to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. Thin provisioning is also called *aggregate overcommitment*.

The storage provided by the aggregate is used up only as reserved LUNs are created or data is appended to files in the volumes.

> **Note:** The aggregate must provide enough free space to hold the metadata for each FlexVol volume it contains. The space required for a FlexVol volume's metadata is approximately 0.5 percent of the volume's configured size.

When the aggregate is overcommitted, it is possible for writes (hole writes or overwrites) to LUNs or files in volumes contained by that aggregate to fail if sufficient free space is not available to accommodate the write.

You can configure a thinly provisioned volume to automatically secure more space from its aggregate when it needs to. However, if you have overcommitted your aggregate, you must monitor your available space and add storage to the aggregate as needed to avoid write errors due to insufficient space.

Aggregates can provide storage to volumes contained by more than one Storage Virtual Machine (SVM). If you are using thin provisioning and you need to maintain strict separation between your SVMs (for example, if you are providing storage in a multi-tenancy environment), you should either use fully allocated volumes (thick provisioning) or ensure that your aggregates are not shared between tenants.

How storage pool works

A *storage pool* is a collection of SSDs. You can combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares across multiple Flash Pool aggregates, at the same time.

Storage pools consist of allocation units, which you can use to provide SSDs and SSD spares to aggregates or to increase the existing SSD size.

After you add an SSD to a storage pool, you can no longer use the SSD as an individual disk. You must use the storage pool to assign or allocate the storage provided by the SSD.

Description of cluster object windows and dialog boxes

You can view all your clusters and cluster objects from the respective storage object page. You can also view the details from the corresponding storage object details page.

Clusters page

The Clusters page enables you to add clusters and to view detailed information about the clusters that you are monitoring.

You must have the OnCommand Administrator or Storage Administrator role.

- *Command buttons* on page 168
- *Clusters list* on page 168
- *Filters pane* on page 170

Command buttons

View Monitoring Status

Enables you to view the monitoring statuses of the selected clusters by navigating to the Manage Data Sources page.

Annotate

Enables you to annotate the selected cluster.

Refresh List

Refreshes the clusters list and the properties associated with the cluster.

Export

Enables you to export the details of all the monitored clusters to a comma-separated values (`.csv`) file.

Clusters table

The Clusters table displays the properties of all the discovered clusters. You can use the column filters to customize the data that is displayed:

Status

Displays an icon that identifies the current status of the cluster. The status can be Critical (✖), Error (!), Warning (⚠), or Normal (✓).

You can position your cursor over the icon to view more information about the event or events generated for the cluster.

If the status of the cluster is based on a single event, you can view information such as the event name, time and date when the event was generated, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event.

If the status of the cluster is based on multiple events of the same severity, the top three events are displayed, along with information such as the event name, time and date when the events are generated, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

Cluster

Displays the name of the cluster.

Communication Status

Displays whether the cluster is reachable or not.

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

System Health

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

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

Host Name or IP Address

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

FQDN

Displays the fully qualified domain name (FQDN) of the cluster. By default, this column is hidden.

OS Version

Displays the Data ONTAP version that the cluster is running.

If the nodes in the cluster are running different versions of Data ONTAP, then the earliest Data ONTAP version is displayed.

Node Count

Displays the number of nodes that belong to the cluster. This column is not sortable.

Last Refreshed Time

Displays the timestamp of when the monitoring samples of the cluster were last collected. This column is not sortable.

Serial Number

Displays the serial number of the cluster.

Contact

Displays the contact information of the cluster. By default, this column is hidden.

Location

Displays the location of the cluster.

FIPS Enabled

Displays whether FIPS mode is enabled on the cluster.

Filters pane

The Filters pane enables you to set filters to customize the display of information in the clusters list. You can select filters in the Status, Communication Status, System Health, and Annotation columns.

> **Note:** The filters specified in the Filters pane override the filters specified for the columns in the clusters list.

Related tasks

Adding clusters on page 20
Editing clusters on page 22
Removing clusters on page 25
Exporting data to CSV files on page 31

Cluster details page

The Cluster details page provides detailed information about a selected cluster that is monitored by Unified Manager, such as health, capacity, and configuration details. You can also view information about the logical interfaces (LIFs), nodes, disks, related devices, and related alerts for the cluster.

- *Command buttons* on page 170
- *Health tab* on page 171
- *Capacity tab* on page 171
- *Configuration tab* on page 173
- *MetroCluster Connectivity tab* on page 175
- *MetroCluster Replication tab* on page 175
- *LIFs tab* on page 175
- *Nodes tab* on page 176
- *Disks tab* on page 178
- *Related Annotations pane* on page 180
- *Related Devices pane* on page 180
- *Related Groups pane* on page 180
- *Related Alerts pane* on page 180

Command buttons

- ★: Enables you to add the selected cluster to the Favorites dashboard of OnCommand Unified Manager.
- **Actions**
    - Add Alert: Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.
    - Rediscover: Initiates a manual refresh of the cluster, which enables Unified Manager to discover recent changes to the cluster.
      If Unified Manager is paired with OnCommand Workflow Automation, the rediscovery operation also reacquires cached data from WFA, if any.

After the rediscovery operation is initiated, a link to the associated job details is displayed to enable tracking of the job status.

  - Annotate: Enables you to annotate the selected cluster.
  - Edit Cluster: Opens the Edit Cluster dialog box, which enables you to edit the settings of the selected cluster.

View Clusters

Enables you to navigate to the Clusters page.

Health tab

Displays detailed information about the data availability and data capacity issues of various cluster objects such as nodes, SVMs, and aggregates. Availability issues are related to the data-serving capability of the cluster objects. Capacity issues are related to the data-storing capability of the cluster objects.

You can click the graph of an object to view a filtered list of the objects. For example, you can click the SVM capacity graph that displays warnings to view a filtered list of SVMs. This list contains SVMs that have volumes or qtrees that have capacity issues with a severity level of Warning. You can also click the SVMs availability graph that displays warnings to view the list of SVMs that have availability issues with a severity level of Warning.

Availability Issues

Graphically displays the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the cluster. For example, information is displayed about disk shelves that are down and aggregates that are offline.

Note: The data displayed for the SFO bar graph is based on the HA state of the nodes. The data displayed for all other bar graphs is calculated based on the events generated.

Capacity Issues

Graphically displays the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the cluster. For example, information is displayed about aggregates that are likely to breach the set threshold values.

Capacity tab

Displays detailed information about the capacity of the selected cluster.

Capacity

Displays details about the used capacity and available capacity from all allocated aggregates:

- Total Capacity

  Displays the total capacity (in MB, GB, and so on) of the cluster. This does not include the capacity that is assigned for parity.
- Used

  Displays the capacity (in MB, GB, and so on) that is used by data. This does not include the capacity that is used for parity, right-sizing, and reservation.
- Available

Displays the capacity (in MB, GB, and so on) available for data.

- Spares
  Displays the storable capacity (in MB, GB, and so on) available for storage in all the spare disks.
- Provisioned
  Displays the capacity (in MB, GB, and so on) that is provisioned for all the underlying volumes.

Capacity Breakout by Disk Type

The Capacity Breakout by Disk Type area displays detailed information about the disk capacity of the various types of disks in the cluster. By clicking the disk type, you can view more information about the disk type from the Disks tab.

- Total Usable Capacity
  Displays the available capacity and spare capacity of the data disks.
- HDD
  Graphically displays the used capacity and available capacity of all the HDD data disks in the cluster. The dotted line represents the spare capacity of the data disks in the HDD.
- Flash
  - SSD Data
    Graphically displays the used capacity and available capacity of the SSD data disks in the cluster.
  - SSD Cache
    Graphically displays the storable capacity of the SSD cache disks in the cluster.
  - SSD Spare
    Graphically displays the spare capacity of the SSD, data, and cache disks in the cluster.
- Unassigned Disks
  Displays the number of unassigned disks in the cluster.

Aggregates with Capacity Issues list

Displays in tabular format details about the used capacity and available capacity of the aggregates that have capacity risk issues.

- Status
  Indicates that the aggregate has a capacity-related issue of a certain severity.

  You can move the pointer over the status to view more information about the event or events generated for the aggregate.

  If the status of the aggregate is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event.

  If the status of the aggregate is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

  Note: An aggregate can have multiple capacity-related events of the same severity or different severities. However, only the highest severity is displayed. For example, if an aggregate has two events with severity levels of Error and Critical, only the Critical severity is displayed.

- Aggregate
  Displays the name of the aggregate.
- Used Data Capacity
  Graphically displays information about the aggregate capacity usage (in percentage).
- Days to Full
  Displays the estimated number of days remaining before the aggregate reaches full capacity.

Configuration tab

Displays details about the selected cluster, such as IP address, serial number, contact, and location:

Cluster Overview

- Management LIF
  Displays the cluster-management LIF that Unified Manager uses to connect to the cluster. The operational status of the LIF is also displayed.
- Host Name or IP Address
  Displays the FQDN, short name, or the IP address of the cluster-management LIF that Unified Manager uses to connect to the cluster.
- FQDN
  Displays the fully qualified domain name (FQDN) of the cluster.
- OS Version
  Displays the Data ONTAP version that the cluster is running. If the nodes in the cluster are running different versions of Data ONTAP, then the earliest Data ONTAP version is displayed.
- Serial Number
  Displays the serial number of the cluster.
- Contact
  Displays details about the administrator whom you should contact in case of issues with the cluster.
- Location
  Displays the location of the cluster.

Remote Cluster Overview

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

- Cluster
  Displays the name of the remote cluster. You can click the cluster name to navigate to the details page of the cluster.
- Hostname or IP Address
  Displays the FQDN, short name, or IP address of the remote cluster.
- Serial Number
  Displays the serial number of the remote cluster.
- Location
  Displays the location of the remote cluster.

MetroCluster Overview

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

- Type
  Displays whether the MetroCluster type is two-node or four-node.
- Configuration
  Displays the MetroCluster configuration, which can have the following values:
  - Stretch Configuration with SAS cables
  - Stretch Configuration with FC-SAS bridge
  - Fabric Configuration with FC switches

  Note: For a four-node MetroCluster, only Fabric Configuration with FC switches is supported.
- Automated Unplanned Switch Over (AUSO)
  Displays whether automated unplanned switchover is enabled for the local cluster. By default, AUSO is enabled for all clusters in a two-node MetroCluster configuration in Unified Manager. You can use the command-line interface to change the AUSO setting.

Nodes

- Availability

  Displays the number of nodes that are up (🟢) or down (🟠) in the cluster.
- OS Versions
  Displays the Data ONTAP versions that the nodes are running as well as the number of nodes running a particular version of Data ONTAP. For example, 8.2 (2), 8.1 (1) specifies that two nodes are running Data ONTAP 8.2, and one node is running Data ONTAP 8.1.

Storage Virtual Machines

- Availability

  Displays the number of SVMs that are up (🟢) or down (🟠) in the cluster.

LIFs

- Availability

  Displays the number of non-data LIFs that are up (🟢) or down (🟠) in the cluster.
- Cluster-Management LIFs
  Displays the number of cluster-management LIFs.
- Node-Management LIFs
  Displays the number of node-management LIFs.
- Cluster LIFs
  Displays the number of cluster LIFs.
- Intercluster LIFs
  Displays the number of intercluster LIFs.

Protocols

- Data Protocols
  Displays the list of licensed data protocols that are enabled for the cluster. The data protocols include iSCSI, CIFS, NFS, and FC and FCoE.

MetroCluster Connectivity tab

Displays the issues and connectivity status of the cluster components in the MetroCluster configuration. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.

> **Note:** The MetroCluster Connectivity tab is displayed only for clusters that are in a MetroCluster configuration.

You can navigate to the details page of a remote cluster by clicking the name of the remote cluster. You can also view the details of the components by clicking the count link of a component. For example, clicking the count link of the node in the cluster displays the node tab in the details page of the cluster. Clicking the count link of the disks in the remote cluster displays the disk tab in the details page of the remote cluster.

You can move the pointer over the components to view the details and the connectivity status of the clusters in case of any issue and to view more information about the event or events generated for the issue.

If the status of the connectivity issue between components is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. The View Details button provides more information about the event.

If status of the connectivity issue between components is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

MetroCluster Replication tab

Displays the status of the data that is being replicated. You can use the MetroCluster Replication tab to ensure data protection by synchronously mirroring the data with the already peered clusters. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.

> **Note:** The MetroCluster Replication tab is displayed only for clusters that are in a MetroCluster configuration.

In a MetroCluster environment, you can use this tab to verify the logical connections and peering of the local cluster with the remote cluster. You can view the objective representation of the cluster components with their logical connections. This helps to identify the issues that might occur during mirroring of metadata and data.

In the MetroCluster Replication tab, local cluster provides the detailed graphical representation of the selected cluster and MetroCluster partner refers to the remote cluster.

LIFs tab

Displays details about all the non-data LIFs that are created on the selected cluster.

LIF

Displays the name of the LIF that is created on the selected cluster.

Operational Status

Displays the operational status of the LIF, which can be Up (⬆), Down (⬇), or Unknown (?). The operational status of a LIF is determined by the status of its physical ports.

Administrative Status

Displays the administrative status of the LIF, which can be Up (⬆), Down (⬇), or Unknown (?). You can control the administrative status of a LIF when you make changes to the configuration or during maintenance. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.

IP Address

Displays the IP address of the LIF.

Role

Displays the role of the LIF. Possible roles are Cluster-Management LIFs, Node-Management LIFs, Cluster LIFs, and Intercluster LIFs.

Home Port

Displays the physical port to which the LIF was originally associated.

Current Port

Displays the physical port to which the LIF is currently associated. After LIF migration, the current port might be different from the home port.

Failover Policy

Displays the failover policy that is configured for the LIF.

Routing Groups

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for clustered Data ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

Failover Group

Displays the name of the failover group.

Nodes tab

Displays information about nodes in the selected cluster. You can view detailed information about the HA pairs, disk shelves, and ports:

HA Details

Provides a pictorial representation of the HA state and the health status of the nodes in the HA pair. The health status of the node is indicated by the following colors:

Green

The node is in a working condition.

Yellow

The node has taken over the partner node or the node is facing some environmental issues.

Red

The node is down.

You can view information about the availability of the HA pair and take required action to prevent any risks. For example, in the case of a possible takeover operation, the following message is displayed: `Storage failover possible`.

You can view a list of the events related to the HA pair and its environment, such as fans, power supplies, NVRAM battery, flash cards, service processor, and connectivity of disk shelves. You can also view the time when the events were triggered.

You can view other node-related information, such as the model number and the serial number.

If there are single-node clusters, you can also view details about the nodes.

Disk Shelves

Displays information about the disk shelves in the HA pair.

You can also view events generated for the disk shelves and the environmental components, and the time when the events were triggered.

Shelf ID

Displays the ID of the shelf where the disk is located.

Component Status

Displays environmental details of the disk shelves, such as power supplies, fans, temperature sensors, current sensors, disk connectivity, and voltage sensors. The environmental details are displayed as icons in the following colors:

Green

The environmental components are in working properly.

Grey

No data is available for the environmental components.

Red

Some of the environmental components are down.

State

Displays the state of the disk shelf. The possible states are Offline, Online, No status, Initialization required, Missing, and Unknown.

Model

Displays the model number of the disk shelf.

Local Disk Shelf

Indicates whether the disk shelf is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

Unique ID

Displays the unique identifier of the disk shelf.

Firmware Version

Displays the firmware version of the disk shelf.

Ports

Displays information about the associated FC, FCoE, and Ethernet ports. You can view details about the ports and the associated LIFs by clicking the port icons.

You can also view the events generated for the ports.

You can view the following port details:

- Port ID
  Displays the name of the port. For example, the port names can be e0M, e0a, and e0b.

- Role
  Displays the role of the port. The possible roles are Cluster, Data, Intercluster, Node-Management, and Undefined.

- Type
  Displays the physical layer protocol used for the port. The possible types are Ethernet, Fibre Channel, and FCoE.

- WWPN
  Displays the World Wide Port Name (WWPN) of the port.

- Firmware Rev
  Displays the firmware revision of the FC/FCoE port.

- Status
  Displays the current state of the port. The possible states are Up, Down, Link Not Connected. or Unknown (?).

You can view the port-related events from the Events list. You can also view the associated LIF details, such as LIF name, operational status, IP address or WWPN, protocols, name of the SVM associated with the LIF, current port, failover policy and failover group.

Disks tab

Displays details about the disks in the selected cluster. You can view disk-related information such as the number of used disks, spare disks, broken disks, and unassigned disks. You can also view other details such as the disk name, disk type, and the owner node of the disk.

Disk Pool Summary

Displays the number of disks, which are categorized by effective types (FCAL, SAS, SATA, MSATA, SSD, Array LUN, and VMDISK), and the state of the disks. You can also view other details, such as the number of aggregate, shared disks, spare disks, broken disks, unassigned disks, and unsupported disks. If you click the effective disk type count link, disks of the selected state and effective type are displayed. For example, if you click the count link for the disk state Broken and effective type SAS, all disks with the disk state Broken and effective type SAS are displayed.

Disk

Displays the name of the disk.

RAID Groups

Displays the name of the RAID group.

Owner Node

Displays the name of the node to which the disk belongs. If the disk is unassigned, no value is displayed in this column.

State

Displays the state of the disk: Aggregate, Shared, Spare, Broken, Unassigned, Unsupported or Unknown. By default, this column is sorted to display the states in the following order: Broken, Unassigned, Unsupported, Spare, Aggregate, and Shared.

Local Disk

Displays either Yes or No to indicate whether the disk is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

Position

Displays the position of the disk based on its container type: for example, Copy, Data, or Parity. By default, this column is hidden.

Impacted Aggregates

Displays the number of aggregates that are impacted due to the failed disk. You can move the pointer over the count link to view the impacted aggregates and then click the aggregate name to view details of the aggregate. You can also click the aggregate count to view the list of impacted aggregates in the Aggregates page.

No value is displayed in this column for the following cases:

- For broken disks when a cluster containing such disks is added to Unified Manager
- When there are no failed disks

Storage Pool

Displays the name of the storage pool to which the SSD belongs. You can move the pointer over the storage pool name to view details of the storage pool.

Storable Capacity

Displays the disk capacity that is available for use.

Raw Capacity

Displays the capacity of the raw, unformatted disk before right-sizing and RAID configuration. By default, this column is hidden.

Type

Displays the types of disks: for example, ATA, SATA, FCAL, or VMDISK.

Effective Type

Displays the disk type assigned by Data ONTAP.

Certain Data ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and spare management. Data ONTAP assigns an effective disk type for each disk type.

Spare Blocks Consumed %

Displays in percentage the spare blocks that are consumed in the SSD disk. This column is blank for disks other than SSD disks.

Rated Life Used %

Displays in percentage an estimate of the SSD life used, based on the actual SSD usage and the manufacturer's prediction of SSD life. A value greater than 99 indicates that the estimated endurance has been consumed, but may not indicate SSD failure. If the value is unknown, then the disk is omitted.

Firmware

Displays the firmware version of the disk.

RPM

Displays the revolutions per minute (RPM) of the disk. By default, this column is hidden.

Model

Displays the model number of the disk. By default, this column is hidden.

Vendor

Displays the name of the disk vendor. By default, this column is hidden.

Shelf ID

Displays the ID of the shelf where the disk is located.

Bay

Displays the ID of the bay where the disk is located.

Related Annotations pane

Enables you to view the annotation details associated with the selected cluster. The details include the annotation name and the annotation values that are applied to the cluster. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

Enables you to view device details that are associated with the selected cluster.

The details include properties of the device that is connected to the cluster such as the device type, size, count, and health status. You can click on the count link for further analysis on that particular device.

You can use MetroCluster Partner pane to obtain count and also details on the remote MetroCluster partner along with its associated cluster components such as nodes, aggregates, and SVMs. The MetroCluster Partner pane is displayed only for clusters in a MetroCluster configuration.

The Related Devices pane enables you to view and navigate to the nodes, SVMs, and aggregates that are related to the cluster:

MetroCluster Partner

> Displays the health status of the MetroCluster partner. Using the count link, you can navigate further and obtain information about the health and capacity of the cluster components.

Nodes

> Displays the number, capacity, and health status of the nodes that belong to the selected cluster. Capacity indicates the total usable capacity over available capacity.

Storage Virtual Machines

> Displays the number of SVMs that belong to the selected cluster.

Aggregates

> Displays the number, capacity, and the health status of the aggregates that belong to the selected cluster.

Related Groups pane

Enables you to view the list of groups that includes the selected cluster.

Related Alerts pane

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

Related tasks

Viewing the cluster list and details on page 147
Adding an alert on page 102
Editing clusters on page 22

Related references

Storage Pool dialog box on page 210

Add Cluster dialog box

You can add an existing cluster to Unified Manager to monitor the cluster and obtain information about its health, capacity, and configuration. You can also associate a Performance Manager instance to the cluster. The cluster is added to both Unified Manager and Performance Manager.

Host Name or IP Address

Specify the FQDN, short name, or the IP address (IPv4 or IPv6) of the cluster-management LIF that is used to connect to the cluster.

User Name

Specify a user name that can be used to log in to the cluster.

Password

Specify a password for the specified user name.

Protocol

Select the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.

Port

Specify the port number of the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Link Performance Manager section

You can associate an instance of Performance Manager to the cluster.

Select Application Instance

Select the instance of Performance Manager that will monitor the cluster.

Command buttons

Save

Adds an existing cluster and closes the Add Cluster dialog box.

Cancel

Discards the changes and closes the Add Cluster dialog box.

Edit Cluster dialog box

The Edit Cluster dialog box enables you to modify the settings of an existing cluster, including the IP address, port, and protocol. For example, you can change the protocol from HTTP to HTTPS. You can also add an instance of Performance Manager to the cluster to monitor the performance.

Host Name or IP Address

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

User Name

Specify a user name that can be used to log in to the cluster.

Password

Specify a password for the specified user name.

Protocol

Select the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.

Port

Specify the port number of the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Link Performance Manager section

You can associate an instance of Performance Manager to the cluster.

Select Application Instance

Select the instance of Performance Manager that will monitor the cluster.

Command buttons

Save

Saves the changes made to the settings of an existing cluster and closes the Edit Cluster dialog box.

Cancel

Discards the changes and closes the Edit Cluster dialog box.

Nodes page

The Nodes page enables you to view detailed information about the nodes in a selected cluster.

- *Command button* on page 182
- *Nodes list* on page 182
- *Filters pane* on page 183

Command button

Export

Enables you to export the details of all the monitored nodes to a comma-separated values (`.csv`) file.

Nodes list

The Nodes list displays the properties of all the discovered nodes in a cluster. You can use the column filters to customize the data that is displayed.

Status

Displays an icon that identifies the current status of the node. The status can be Critical (✖), Error (!), Warning (⚠), or Normal (✓).

You can position your cursor over the icon to view more information about the event or events generated for the node.

Node

Displays the name of the node.

State

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

HA State

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

Down Time

Displays the time that has elapsed or the timestamp since the node is offline. If the time elapsed exceeds a week, the timestamp when the node went offline is displayed. By default, this column is hidden.

Cluster

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

Model

Displays the model of the node.

OS version

Displays the ONTAP software version that the node is running.

All Flash Optimized

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

Serial Number

Displays the serial number of the node. By default, this column is hidden.

Firmware Version

Displays the firmware version number of the node. By default, this column is hidden.

Owner

Displays the name of the node's owner. By default, this column is hidden.

Location

Displays the location of the node.

Aggregate Used Capacity

Displays the amount of space (in GB or MB) used for data in the node's aggregates.

Aggregate Total Capacity

Displays the total space (in GB or MB) available for data in the node's aggregates.

Usable Spare Capacity

Displays the amount of available space (in GB or MB) in the node that can be used to enhance the aggregate capacity.

Usable Raw Capacity

Displays the amount of space (in GB or MB) that is usable in the node.

Total Raw Capacity

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

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the nodes list. You can select filters related to the Status, State, and HA State columns.

Note: The filters that are specified in the Filters pane override the filters that are specified for the columns in the Nodes list.

Related tasks

Exporting data to CSV files on page 31

Storage Virtual Machines page

The Storage Virtual Machines page enables you to view detailed information about the Storage Virtual Machines (SVMs) that you are monitoring.

- *Command button* on page 184
- *SVMs list* on page 184
- *Filters pane* on page 185

Command buttons

Export

Enables you to export the details of all the monitored SVMs to a comma-separated values (`.csv`) file.

Annotate

Enables you to annotate the selected Storage Virtual Machine (SVM).

SVMs list

The SVMs list displays, in tabular format, the properties of all the discovered SVMs. You can use the column filters to customize the data that is displayed:

Status

Displays the current status of the SVM. The status can be Critical (✖), Error (!), Warning (⚠), or Normal (✓).

You can move the pointer over the status to view more information about the event or events generated for the SVM.

If the status of the SVM is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the View Details button to view more information about the event.

If the status of the SVM is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the View All Events link to view the list of generated events.

Storage Virtual Machine

Displays the name of the SVM.

You can move the pointer over each SVM to view information such as the last generated event, cluster to which the SVM belongs, volume type of the SVM, allowed protocols, and space allocated in the SVM. You can also view the details of related objects such as the cluster to which the SVM belongs, all the SVMs that belong to the cluster, and the volumes that belong to the SVM.

State

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

Cluster

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

Allowed Volume Type

Displays the type of volume that can be created in the SVM. The type can be InfiniteVol or FlexVol.

Available Data Capacity

Displays the available data capacity (in MB, GB, and so on) of all the volumes in the SVM.

Total Data Capacity

Displays the total data capacity (in MB, GB, and so on) of all the volumes in the SVM.

Root Volume

Displays the name of the root volume of the SVM. By default, this column is hidden.

NIS State

Displays the state of the Network Information Service (NIS). The state can be Enabled, Disabled, or Not Configured. By default, this column is hidden.

NIS Domain

Displays the NIS domain name. This column is blank when the NIS server is disabled or is not configured. By default, this column is hidden.

DNS State

Displays the state of the Domain Name System (DNS). The state can be Enabled, Disabled, or Not Configured. By default, this column is hidden.

DNS Domain

Displays the DNS domain name. By default, this column is hidden.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the SVMs list. You can select filters related to the Status, State, and Annotation columns.

Note: The filters specified in the Filters pane override the filters specified for the columns in the SVMs list.

Related tasks

Viewing the details of SVMs with Infinite Volume on page 360
Exporting data to CSV files on page 31

Storage Virtual Machine details page

You can use the Storage Virtual Machine details page to view detailed information about the selected Storage Virtual Machine (SVM, formerly known as Vserver) that is monitored by Unified Manager, such as its health, capacity, configuration, data policies, logical interfaces (LIFs), LUNs, qtrees, and user and user group quotas. You can also view information about the related objects and related alerts for the SVM.

Note: You can monitor only data SVMs.

- *Command buttons* on page 186
- *Health tab* on page 186
- *Capacity tab* on page 187
- *Configuration tab* on page 189
- *LIFs tab* on page 190
- *Qtrees tab* on page 191

- *User and Group Quotas tab* on page 193
- *NFS Exports tab* on page 194
- *CIFS Shares tab* on page 195
- *SAN tab* on page 197
- *Data Policy tab* on page 197
- *Related Annotations pane* on page 199
- *Related Devices pane* on page 199
- *Related Groups pane* on page 199
- *Related Alerts pane* on page 199

Command buttons

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

Actions

- Add Alert
  Enables you to add an alert to the selected SVM.
- Edit Thresholds
  Enables you to edit the SVM thresholds.

  Note: This button is enabled only for SVM with Infinite Volume.
- Annotate
  Enables you to annotate the selected SVM.

View Storage Virtual Machines

Enables you to navigate to the Storage Virtual Machines page.

Health tab

The Health tab displays detailed information about data availability, data capacity, and protection issues of various objects such as volumes, aggregates, NAS LIFs, SAN LIFs, LUNs, protocols, services, NFS exports, and CIFS shares.

You can click the graph of an object to view the filtered list of objects. For example, you can click the volume capacity graph that displays warnings to view the list of volumes that have capacity issues with severity as warning.

Availability Issues

Displays, as a graph, the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the SVM. For example, information is displayed about the NAS LIFs and the SAN LIFs that are down and volumes that are offline.

You can also view information about the related protocols and services that are currently running, and the number and status of NFS exports and CIFS shares.

If the selected SVM is an SVM with Infinite Volume, you can view availability details about the Infinite Volume.

Capacity Issues

Displays, as a graph, the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the SVM. For example, information is displayed about aggregates that are likely to breach the set threshold values.

If the selected SVM is an SVM with Infinite Volume, you can view capacity details about the Infinite Volume.

Protection Issues

Provides a quick overview of SVM protection-related health by displaying, as a graph, the total number of relationships, including relationships that have protection issues and relationships that do not have any protection-related issues. When unprotected volumes exist, clicking on the link takes you to the Volumes page where you can view a filtered list of the unprotected volumes on the SVM. The colors in the graph represent the different severity levels of the issues. Clicking a graph takes you to the Volume Protection Relationships page, where you can view a filtered list of protection relationship details. The information below the graph provides details about protection issues that can impact or have already impacted the protection of data in the SVM. For example, information is displayed about volumes that have a Snapshot copy reserve that is almost full or about SnapMirror relationship lag issues.

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

Capacity tab

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

The following information is displayed for an SVM with FlexVol volume:

Capacity

The Capacity area displays details about the used and available capacity allocated from all volumes:

- Total Capacity
  Displays the total capacity (in MB, GB, and so on) of the SVM.
- Used
  Displays the space used by data in the volumes that belong to the SVM.
- Guaranteed Available
  Displays the guaranteed available space for data that is available for volumes in the SVM.
- Unguaranteed
  Displays the available space remaining for data that is allocated for thinly provisioned volumes in the SVM.

Volumes with Capacity Issues

The Volumes with Capacity Issues list displays, in tabular format, details about the volumes that have capacity issues:

- Status
  Indicates that the volume has a capacity-related issue of a certain severity.
  You can move the pointer over the status to view more information about the capacity-related event or events generated for the volume.
  If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use the **View Details** button to view more information about the event.

If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

> **Note:** A volume can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a volume has two events with severities of Error and Warning, only the Error severity is displayed.

- Volume
  Displays the name of the volume.
- Used Data Capacity
  Displays, as a graph, information about the volume capacity usage (in percentage).
- Days to Full
  Displays the estimated number of days remaining before the volume reaches full capacity.
- Thin Provisioned
  Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.
- Aggregate
  Displays the name of the aggregate that contains the volume.

The following information is displayed for an SVM with Infinite volume:

Capacity

Displays the following capacity-related details:

- Percentage of used and free data capacity
- Percentage of used and free Snapshot capacity
- Snapshot Overflow
  Displays the data space that is consumed by the Snapshot copies.
- Used
  Displays the space used by data in the SVM with Infinite Volume.
- Warning
  Indicates that the space in the SVM with Infinite Volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.
- Error
  Indicates that the space in the SVM with Infinite Volume if full. If this threshold is breached, the Space Full event is generated.

Other Details

- Total Capacity
  Displays the total capacity in the SVM with Infinite Volume.
- Data Capacity
  Displays used data capacity, available data capacity, and Snapshot overflow capacity details of the SVM with Infinite Volume.
- Snapshot Reserve
  Displays the used and free details of the Snapshot reserve.

- System Capacity
  Displays the used system capacity and available system capacity in the SVM with Infinite Volume.
- Thresholds
  Displays the nearly full and full thresholds of the SVM with Infinite Volume.

Storage Class Capacity Details

Displays information about the capacity usage in your storage classes. This information is displayed only if you have configured storage classes for your SVM with Infinite Volume.

Storage Virtual Machine Storage Class Thresholds

Displays the following thresholds (in percentage) of your storage classes:

- Nearly Full Threshold
  Specifies the percentage at which a storage class in an SVM with Infinite Volume is considered to be nearly full.
- Full Threshold
  Specifies the percentage at which the storage class in an SVM with Infinite Volume is considered full.
- Snapshot Usage Limit
  Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.

Configuration tab

The Configuration tab displays configuration details about the selected SVM, such as its cluster, root volume, the type of volumes it contains (Infinite Volume or FlexVol volumes), and the policies created on the SVM:

Overview

- Cluster
  Displays the name of the cluster to which the SVM belongs.
- Allowed Volume Type
  Displays the type of volumes that can be created in the SVM. The type can be InfiniteVol or FlexVol.
- Root Volume
  Displays the name of the root volume of the SVM.
- Allowed Protocols
  Displays the type of protocols that can be configured on the SVM. Also, indicates if a protocol is up (🟢), down (🟠), or is not configured (⚪).

Data LIFs

- NAS
  Displays the number of NAS LIFs that are associated with the SVM. Also, indicates if the LIFs are up (🟢) or down (🟠).
- SAN
  Displays the number of SAN LIFs that are associated with the SVM. Also, indicates if the LIFs are up (🟢) or down (🟠).
- Junction Path

Displays the path on which the Infinite Volume is mounted. Junction path is displayed for an SVM with Infinite Volume only.

- Storage Classes
  Displays the storage classes associated with the selected SVM with Infinite Volume. Storage classes are displayed for an SVM with Infinite Volume only.

Management LIFs

- Availability
  Displays the number of management LIFs that are associated with the SVM. Also, indicates if the management LIFs are up (🟢) or down (🟠).

Policies

- Snapshots
  Displays the name of the Snapshot policy that is created on the SVM.
- Export Policies
  Displays either the name of the export policy if a single policy is created or displays the number of export policies if multiple policies are created.
- Data Policy
  Displays whether a data policy is configured for the selected SVM with Infinite Volume.

Services

- Type
  Displays the type of service that is configured on the SVM. The type can be Domain Name System (DNS) or Network Information Service (NIS).
- State
  Displays the state of the service, which can be Up (🟢), Down (🟠), or Not Configured (⚪).
- Domain Name
  Displays the fully qualified domain names (FQDNs) of the DNS server for the DNS services or NIS server for the NIS services. When the NIS server is enabled, the active FQDN of the NIS server is displayed. When the NIS server is disabled, the list of all the FQDNs are displayed.
- IP Address
  Displays the IP addresses of the DNS or NIS server. When the NIS server is enabled, the active IP address of the NIS server is displayed. When the NIS server is disabled, the list of all the IP addresses are displayed.

LIFs tab

The LIFs tab displays details about the data LIFs that are created on the selected SVM:

LIF

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

Operational Status

Displays the operational status of the LIF, which can be Up (⬆), Down (⬇), or Unknown (❓). The operational status of a LIF is determined by the status of its physical ports.

Administrative Status

Displays the administrative status of the LIF, which can be Up (⬆), Down (⬇), or Unknown (?). The administrative status of a LIF is controlled by the storage administrator to make changes to the configuration or for maintenance purposes. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.

IP Address / WWPN

Displays the IP address for Ethernet LIFs and the World Wide Port Name (WWPN) for FC LIFs.

Protocols

Displays the list of data protocols that are specified for the LIF, such as CIFS, NFS, iSCSI, FC/FCoE, and FlexCache. For Infinite Volume, the SAN protocols are not applicable.

Role

Displays the LIF role. The roles can be Data or Management.

Home Port

Displays the physical port to which the LIF was originally associated.

Current Port

Displays the physical port to which the LIF is currently associated. If the LIF is migrated, the current port might be different from the home port.

Port Set

Displays the port set to which the LIF is mapped.

Failover Policy

Displays the failover policy that is configured for the LIF. For NFS, CIFS, and FlexCache LIFs, the default failover policy is Next Available. Failover policy is not applicable for FC and iSCSI LIFs.

Routing Groups

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for clustered Data ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

Failover Group

Displays the name of the failover group.

Qtrees tab

The Qtrees tab displays details about qtrees and their quotas.

Note: The Qtrees tab is not displayed for an SVM with Infinite Volume.

Status

Displays the current status of the qtree. The status can be Critical (✖), Error (!), Warning (⚠), or Normal (✓).

You can move the pointer over the status icon to view more information about the event or events generated for the qtree.

If the status of the qtree is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the qtree is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

> **Note:** A qtree can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a qtree has two events with severities of Error and Warning, only the Error severity is displayed.

Qtree

Displays the name of the qtree.

Volume

Displays the name of the volume that contains the qtree.

You can move the pointer over the volume name to view more information about the volume.

Quota Set

Indicates whether a quota is enabled or disabled on the qtree.

Disk Used %

Displays the percentage of disk space used. If a disk hard limit is set, this value is based on the disk hard limit. If the quota is set without a disk hard limit, the value is based on the volume data space. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then “Not applicable” is displayed.

Disk Hard Limit

Displays the maximum amount of disk space allocated for the qtree. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

Disk Soft Limit

Displays the amount of disk space allocated for the qtree before a warning event is generated. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

Disk Threshold

Displays the threshold value set on the disk space. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk threshold limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

Files Used %

Displays the percentage of files used in the qtree. If the file hard limit is set, this value is based on the file hard limit. No value is displayed if the quota is set without a file hard limit. The value is displayed as “Not applicable” if the quota is not set or if quotas are off on the volume to which qtree belongs.

File Hard Limit

Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

File Soft Limit

Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

User and Group Quotas tab

Displays details about the user and user group quotas for the selected SVM. You can view information such as the status of the quota, name of the user or user group, soft and hard limits set on the disks and files, amount of disk space and number of files used, and the disk threshold value. You can also change the email address associated with a user or user group.

Edit Email Address

Opens the Edit Email Address dialog box, which displays the current email address of the selected user or user group. You can modify the email address. If the **Edit Email Address** field is blank, the default rule is used to generate an email address for the selected user or user group.

If more than one user has the same quota, the names of the users are displayed as comma-separated values. Also, the default rule is not used to generate the email address; therefore, you must provide the required email address for notifications to be sent.

Configure Email Rules

Enables you to create or modify rules to generate an email address for the user or user group quotas that are configured on the SVM. A notification is sent to the specified email address when there is a quota breach.

Status

Displays the current status of the quota. The status can be Critical (✖), Warning (⚠), or Normal (✓).

You can move the pointer over the status icon to view more information about the event or events generated for the quota.

If the status of the quota is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the quota is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

Note: A quota can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a quota has two events with severities of Error and Warning, only the Error severity is displayed.

User or Group

Displays the name of the user or user group. If more than one user has the same quota, the names of the users are displayed as comma-separated values.

The value is displayed as “Unknown” when Data ONTAP does not provide a valid user name because of SecD errors.

Type

Specifies if the quota is for a user or a user group.

Volume or Qtree

Displays the name of the volume or qtree on which the user or user group quota is specified.

You can move the pointer over the name of the volume or qtree to view more information about the volume or qtree.

Disk Used %

Displays the percentage of disk space used. The value is displayed as “Not applicable” if the quota is set without a disk hard limit.

Disk Hard Limit

Displays the maximum amount of disk space allocated for the quota. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” if the quota is set without a disk hard limit.

Disk Soft Limit

Displays the amount of disk space allocated for the quota before a warning event is generated. The value is displayed as “Unlimited” if the quota is set without a disk soft limit. By default, this column is hidden.

Disk Threshold

Displays the threshold value set on the disk space. The value is displayed as “Unlimited” if the quota is set without a disk threshold limit. By default, this column is hidden.

Files Used %

Displays the percentage of files used in the qtree. The value is displayed as “Not applicable” if the quota is set without a file hard limit.

File Hard Limit

Displays the hard limit for the number of files permitted on the quota. The value is displayed as “Unlimited” if the quota is set without a file hard limit.

File Soft Limit

Displays the soft limit for the number of files permitted on the quota. The value is displayed as “Unlimited” if the quota is set without a file soft limit. By default, this column is hidden.

Email Address

Displays the email address of the user or user group to which notifications are sent when there is a breach in the quotas.

NFS Exports tab

The NFS Exports tab displays information about NFS exports such as its status, the path associated with the volume (Infinite Volumes or FlexVol volumes), access levels of clients to the NFS exports, and the export policy defined for the volumes that are exported. NFS exports will not be displayed in the following conditions: if the volume is not mounted or if the protocols associated with the export policy for the volume do not contain NFS exports.

Status

Displays the current status of the NFS export. The status can be Error (❗) or Normal (✓).

Junction Path

Displays the path to which the volume is mounted. If an explicit NFS exports policy is applied to a qtree, the column displays the path of the volume through which the qtree can be accessed.

Junction Path Active

Displays whether the path to access the mounted volume is active or inactive. In Data ONTAP 8.2 and earlier, the status column is green for a root volume and the Junction Path Active column is blank.

Volume or Qtree

Displays the name of the volume or qtree to which the NFS export policy is applied. For Infinite Volumes, the name of the SVM with the Infinite Volume is displayed. If an NFS export policy is applied to a qtree in the volume, the column displays both the names of the volume and the qtree.

You can click the link to view details about the object in the respective details page. If the object is a qtree, links are displayed for both the qtree and the volume.

Volume State

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

- Offline
  Read or write access to the volume is not allowed.
- Online
  Read and write access to the volume is allowed.
- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Security Style

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

- UNIX (NFS clients)
  Files and directories in the volume have UNIX permissions.
- Unified
  Files and directories in the volume have a unified security style.
- NTFS (CIFS clients)
  Files and directories in the volume have Windows NTFS permissions.
- Mixed
  Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

UNIX Permission

Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.

Export Policy

Displays the rules that define the access permission for volumes that are exported. You can click the link to view details about the rules associated with the export policy such as the authentication protocols and the access permission.

CIFS Shares tab

Displays information about the CIFS shares on the selected SVM. You can view information such as the status of the CIFS share, share name, path associated with the SVM, the status of the junction path of the share, containing object, state of the containing volume, security data of the share, and export policies defined for the share. You can also determine whether an equivalent NFS path for the CIFS share exists.

View User Mapping

Launches the User Mapping dialog box.

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

Show ACL

Launches the Access Control dialog box for the share.

You can view user and permission details for the selected share.

Status

Displays the current status of the share. The status can be Normal (✓) or Error (❗).

Share Name

Displays the name of the CIFS share.

Path

Displays the junction path on which the share is created.

Junction Path Active

Displays whether the path to access the share is active or inactive. In Data ONTAP version 8.2 and below, for a root volume the status column is green and the Junction Path Active column is blank.

Containing Object

Displays the name of the containing object to which the share belongs. The containing object can be a volume or a qtree.

By clicking the link, you can view details about the containing object in the respective Details page. If the containing object is a qtree, links are displayed for both qtree and volume.

Volume State

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

- Offline
  Read or write access to the volume is not allowed.
- Online
  Read and write access to the volume is allowed.
- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Security

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

- UNIX (NFS clients)
  Files and directories in the volume have UNIX permissions.
- NTFS (CIFS clients)
  Files and directories in the volume have Windows NTFS permissions.
- Mixed
  Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

Export Policy

Displays the name of the export policy applicable to the share. If an export policy is not specified for the SVM, the value is displayed as Not Enabled.

You can click the link to view details about the rules associated with the export policy, such as access protocols and permissions. The link is disabled if the export policy is disabled for the selected SVM.

NFS Equivalent

Specifies whether there is an NFS equivalent for the share.

SAN tab

Displays details about LUNs, initiator groups, and initiators for the selected SVM. By default, the LUNs view is displayed. You can view details about the initiator groups in the Initiator Groups tab and details about initiators in the Initiators tab.

LUNs

Displays details about the LUNs that belong to the selected SVM. You can view information such as the LUN name, LUN state (online or offline), the name of the file system (volume or qtree) that contains the LUN, the type of host operating system, the total data capacity and serial number of the LUN. You can also view information whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group.

You can also view the initiator groups and initiators that are mapped to the selected LUN.

Initiator Groups

Displays details about initiator groups. You can view details such as the name of the initiator group, the access state, the type of host operating system that is used by all the initiators in the group, and the supported protocol. When you click the link in the access state column, you can view the current access state of the initiator group.

Normal

The initiator group is connected to multiple access paths.

Single Path

The initiator group is connected to a single access path.

No Paths

There is no access path connected to the initiator group.

You can view whether initiator groups are mapped to all the LIFs or specific LIFs through a port set. When you click the count link in the Mapped LIFs column, either all LIFs are displayed or specific LIFs for a port set are displayed. LIFs that are mapped through the target portal are not displayed. The total number of initiators and LUNs that are mapped to an initiator group is displayed.

You can also view the LUNs and initiators that are mapped to the selected initiator group.

Initiators

Displays the name and type of the initiator and the total number of initiator groups mapped to this initiator for the selected SVM.

You can also view the LUNs and initiator groups that are mapped to the selected initiator group.

Data Policy tab

The Data Policy tab enables you to create, modify, activate, or delete one or more rules in a data policy. You can also import the data policy into the Unified Manager database and export the data policy to your computer:

Note: The Data Policy tab is displayed only for SVMs with Infinite Volume.

Rules list

Displays the list of rules. By expanding the rule, you can view the corresponding matching criteria of the rule and the storage class where the content is placed based on the rule.

The default rule is the last rule in the list. You cannot change the order of the default rule.

- Matching Criteria
  Displays the conditions for the rule. For example, a rule can be "File path starts with `/eng/nightly`".

  Note: The file path must always start with a junction path.

- Content Placement
  Displays the corresponding storage class for the rule.

Rule Filter

Enables you to filter rules associated with a specific storage class listed in the list.

Action buttons

- Create
  Opens the Create Rule dialog box, which enables you to create a new rule for your data policy.
- Edit
  Opens the Edit Rule dialog box, which enables you to modify rule properties such as directory paths, file types, and owners.
- Delete
  Deletes the selected rule.
- Move Up
  Moves the selected rule up in the list. However, you cannot move the default rule up in the list.
- Move Down
  Moves the selected rule down the list. However, you cannot move the default rule down the list.
- Activate
  Activates the rules and changes made to the data policy in the SVM with Infinite Volume.
- Reset
  Resets all changes made to the data policy configuration.
- Import
  Imports a data policy configuration from a file.
- Export
  Exports a data policy configuration to a file.

Related Devices area

The Related Devices area enables you to view and navigate to the LUNs, CIFS shares, and the user and user group quotas that are related to the qtree:

LUNs

Displays the total number of the LUNs associated with the selected qtree.

NFS exports

Displays the total number of NFS export policies associated with the selected qtree.

CIFS Shares

Displays the total number of CIFS shares associated with the selected qtree.

User and Group Quotas

Displays the total number of the user and user group quotas associated with the selected qtree. The health status of the user and user group quotas is also displayed, based on the highest severity level.

Related Annotations pane

The Related Annotations pane enables you to view the annotation details associated with the selected SVM. Details include the annotation name and the annotation values that are applied to the SVM. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view the cluster, aggregates, and volumes that are related to the SVM:

Cluster

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

Aggregates

Displays the number of aggregates that belong to the selected SVM. The health status of the aggregates is also displayed, based on the highest severity level. For example, if an SVM contains ten aggregates, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.

Assigned Aggregates

Displays the number of aggregates that are assigned to an SVM. The health status of the aggregates is also displayed, based on the highest severity level.

Volumes

Displays the number and capacity of the volumes that belong to the selected SVM. The health status of the volumes is also displayed, based on the highest severity level.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected SVM.

Related Alerts pane

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

Related tasks

Viewing the details of SVMs with Infinite Volume on page 360
Viewing the constituents of an Infinite Volume on page 360
Editing the Infinite Volume threshold settings on page 361
Creating rules on page 364
Viewing rules on page 366
Previewing changes to your data policy on page 369
Adding an alert on page 102

Related references

Export Policy Rules dialog box on page 229

Aggregates page

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

- *Command buttons* on page 200
- *Aggregates list* on page 200
- *Filters pane* on page 203

Command buttons

Edit Thresholds

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

Export

Enables you to export the details of all the monitored aggregates to a comma-separated values (`.csv`) file.

Aggregates list

Displays, in tabular format, the properties of all the discovered aggregates. You can use the column filters to customize the data that is displayed:

Status

Displays the current status of the aggregate. The status can be Critical (✖), Error (!), Warning (⚠), or Normal (✓).

You can move the pointer over the status to view more information about the event or events generated for the aggregate.

If the status of the aggregate is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event.

If the status of the aggregate is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

Aggregate

Displays the name of the aggregate.

You can move the pointer over an aggregate to view information such as the last generated event, node that contains the aggregate, RAID type, Snapshot reserve, Snapshot copies, and space allocated in the aggregate. You can also view the number of volume move operations that are currently in progress.

State

Displays the current state of the aggregate, which can be one of the following:

- Offline
  Read or write access is not allowed.

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

Node

Displays the name of the storage controller that contains the aggregate.

Mirror Status

Displays the state of the aggregate, which can be one of the following:

- Mirrored
  The aggregate plex data is mirrored.
- Mirror degraded
  The aggregate plex data cannot be mirrored.

- Mirror resynchronizing
  The aggregate plex data is being mirrored.
- Failed
  The aggregate plex data mirroring failed.
- Invalid configuration
  The initial state before an aggregate is created.
- Uninitialized
  The aggregate is being created.
- Unmirrored
  The aggregate is not mirrored.
- CP count check in progress
  The aggregate has been assimilated and Unified Manager is validating that the CP counts for the plexes is similar.
- Limbo
  There is an issue with the aggregate labels. The Data ONTAP system identifies the aggregate but cannot accurately assimilate the aggregate.
- Needs CP count check
  The aggregate is assimilated but the CP counts on both plexes are not yet validated to be similar.

When an aggregate is in the mirror_resynchronizing state, then the resynchronization percentage is also shown. By default, this column is hidden.

In Transition

Indicates whether the aggregate has completed transition or not. This column is hidden by default.

Type

Indicates whether the aggregate is a Flash Pool aggregate (combines HDDs and SSDs), or whether the disks in the aggregate are standard disks (HDDs only) or SSD disks (SSDs only).

For standard disks and SSD disks, this column is blank when the monitored storage system is running a version of clustered Data ONTAP earlier than 8.3.

SnapLock Type

Displays the aggregate SnapLock Type, the available options are Compliance, Enterprise, Non-SnapLock.

Used Data Capacity

Displays the amount of space (in MB, GB, and so on) used for data in the aggregate. By default, this column is hidden.

Used Data %

Displays the percentage of space used for data in the aggregate. By default, this column is hidden.

Available Data Capacity

Displays the amount of space (in MB, GB, and so on) available for data in the aggregate.

Available Data %

Displays the percentage of space available for data in the aggregate. By default, this column is hidden.

Total Data Capacity

Displays the total data size (in MB, GB, and so on) of the aggregate.

Committed Capacity

Displays the total space (in MB, GB, and so on) committed for all of the volumes in the aggregate.

RAID Type

Displays the RAID configuration type, which can be one of the following:

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

Filters pane

Enables you to set filters to customize the way information is displayed in the aggregates list. You can select filters related to the Status column.

Note: The filters specified in the Filters pane override the filters specified for the columns in the aggregates list.

Related tasks

Editing aggregate threshold settings on page 128
Exporting data to CSV files on page 31

Aggregate details page

You can use the Aggregate details page to view detailed information about the selected aggregate that is monitored by Unified Manager, such as the capacity, disk information, configuration details, and events generated. You can also view information about the related objects and related alerts for that aggregate.

- *Command buttons* on page 203
- *Capacity tab* on page 204
- *Disk Information tab* on page 206
- *Configuration tab* on page 208
- *History area* on page 209
- *Events list* on page 209
- *Related Devices pane* on page 209
- *Related Alerts pane* on page 210

Command buttons

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

- ★: Enables you to add the selected aggregate to the Favorites dashboard of OnCommand Unified Manager.

Actions

- Add Alert
  Enables you to add an alert to the selected aggregate.
- Edit Thresholds
  Enables you to modify the threshold settings for the selected aggregate.

View Aggregates

Enables you to navigate to the Aggregates page.

Capacity tab

The Capacity tab displays detailed information about the selected aggregate, such as its capacity, thresholds, and daily growth rate. By default, capacity events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are modified by a technical support representative, the threshold values are applied to the node root aggregate.

Capacity

Displays the data capacity graph and the Snapshot copies graph, which display capacity details about the aggregate:

- Snapshot Overflow
  Displays the data space that is consumed by the Snapshot copies.
- Used
  Displays the space used by data in the aggregate.
- Overcommitted
  Indicates that the space in the aggregate is overcommitted.
- Warning
  Indicates that the space in the aggregate is nearly full. If this threshold is breached, the Space Nearly Full event is generated.
- Error
  Indicates that the space in the aggregate is full. If this threshold is breached, the Space Full event is generated.
- Data graph
  Displays the total data capacity and the used data capacity of the aggregate. If the aggregate is overcommitted, a flag is displayed with the overcommitted capacity.
- Snapshot Copies graph
  This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both of the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

> **Note:** Unified Manager does not display snapshot overflow for systems running Data ONTAP 8.2 or later. Snapshot Autodelete option is enabled in systems running Data ONTAP 8.2 or later, which does not cause snapshot overflow.

Total Capacity

Displays the total capacity in the aggregate.

Data Capacity

Displays the amount of space used by the aggregate (used capacity) and the amount of available space in the aggregate (free capacity).

Snapshot Reserve

Displays the used and free Snapshot capacity of the aggregate.

Overcommitted Capacity

Displays the aggregate overcommitment. Aggregate overcommitment enables you to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. When thin provisioning is in use, the total size of volumes in the aggregate can exceed the total capacity of the aggregate.

Note: If you have overcommitted your aggregate, you must monitor its available space carefully and add storage as required to avoid write errors due to insufficient space.

Total Cache Space

Displays the total space of the solid-state drives (SSDs) or allocation units that are added to a Flash Pool aggregate. If you have enabled Flash Pool for an aggregate but have not added any SSDs, then the cache space is displayed as 0 KB.

Note: This field is hidden if Flash Pool is disabled for an aggregate.

Aggregate Thresholds

Displays the following aggregate capacity thresholds:

- Nearly Full Threshold
  Specifies the percentage at which an aggregate is nearly full.
- Full Threshold
  Specifies the percentage at which an aggregate is full.
- Nearly Overcommitted Threshold
  Specifies the percentage at which an aggregate is nearly overcommitted.
- Overcommitted Threshold
  Specifies the percentage at which an aggregate is overcommitted.

Daily Growth Rate

Displays the disk space used in the aggregate if the rate of change between the last two samples continues for 24 hours.

For example, if an aggregate uses 10 GB of disk space at 2 pm and 12 GB at 6 pm, the daily growth rate (GB) for this aggregate is 2 GB.

Volume Move

Displays the number of volume move operations that are currently in progress:

- Volumes Out
  Displays the number and capacity of the volumes that are being moved out of the aggregate.
  You can click the link to view more details, such as the volume name, aggregate to which the volume is moved, status of the volume move operation, and the estimated end time.
- Volumes In
  Displays the number and remaining capacity of the volumes that are being moved into the aggregate.
  You can click the link to view more details, such as the volume name, aggregate from which the volume is moved, status of the volume move operation, and the estimated end time.

- Estimated used capacity after volume move
  Displays the estimated amount of used space (as a percentage, and in KB, MB, GB, and so on) in the aggregate after the volume move operations are complete.

Capacity Overview - Volumes

Displays graphs that provide information about the capacity of the volumes contained in the aggregate. The amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume is displayed. When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

You can select the graph you want to view from the drop-down lists. You can sort the data displayed in the graph to display details such as the used size, provisioned size, available capacity, fastest daily growth rate, and slowest growth rate. You can filter the data based on the Storage Virtual Machines (SVMs) that contain the volumes in the aggregate. You can also view details for thinly provisioned volumes. You can view the details of specific points on the graph by positioning your cursor over the area of interest. By default, the graph displays the top 30 filtered volumes in the aggregate.

Disk Information tab

Displays detailed information about the disks in the selected aggregate, including the RAID type and size, and the type of disks used in the aggregate. The tab also graphically displays the RAID groups, and the types of disks used (such as SAS, ATA, FCAL, SSD, or VMDISK). You can view more information, such as the disk's bay, shelf, and rotational speed, by positioning your cursor over the parity disks and data disks.

Data

Graphically displays details about dedicated data disks, shared data disks, or both. When the data disks contain shared disks, graphical details of the shared disks are displayed. When the data disks contain dedicated disks and shared disks, graphical details of both the dedicated data disks and the shared data disks are displayed.

RAID Details

RAID details are displayed only for dedicated disks.

- Type
  Displays the RAID type (RAID0, RAID4, RAID-DP, or RAID-TEC).
- Group Size
  Displays the maximum number of disks allowed in the RAID group.
- Groups
  Displays the number of RAID groups in the aggregate.

Disks Used

- Effective Type
  Displays the types of data disks (for example, ATA, SATA, FCAL, SSD, or VMDISK) in the aggregate.
- Data Disks
  Displays the number and capacity of the data disks that are assigned to an aggregate. Data disk details are not displayed when the aggregate contains only shared disks.
- Parity Disks

Displays the number and capacity of the parity disks that are assigned to an aggregate. Parity disk details are not displayed when the aggregate contains only shared disks.

- Shared Disks
  Displays the number and capacity of the shared data disks that are assigned to an aggregate. Shared disk details are displayed only when the aggregate contains shared disks.

Spare Disks

Displays the disk effective type, number, and capacity of the spare data disks that are available for the node in the selected aggregate.

Note: When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

SSD Cache

Provides details about dedicated cache SSD disks and shared cache SSD disks.

The following details for the dedicated cache SSD disks are displayed:

RAID Details

- Type
  Displays the RAID type (RAID0, RAID4, RAID-DP or RAID-TEC).
- Group Size
  Displays the maximum number of disks allowed in the RAID group.
- Groups
  Displays the number of RAID groups in the aggregate.

Disks Used

- Effective Type
  Indicates that the disks used for cache in the aggregate are of type SSD.
- Data Disks
  Displays the number and capacity of the data disks that are assigned to an aggregate for cache.
- Parity Disks
  Displays the number and capacity of the parity disks that are assigned to an aggregate for cache.

Spare Disks

Displays the disk effective type, number, and capacity of the spare disks that are available for the node in the selected aggregate for cache.

Note: When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

Provides the following details for the shared cache:

Storage Pool

Displays the name of the storage pool. You can move the pointer over the storage pool name to view the following details:

- Status
  Displays the status of the storage pool, which can be healthy or unhealthy.
- Total Allocations
  Displays the total allocation units and the size in the storage pool.

- Allocation Unit Size
  Displays the minimum amount of space in the storage pool that can be allocated to an aggregate.
- Disks
  Displays the number of disks used to create the storage pool. If the disk count in the storage pool column and the number of disks displayed in the Disk Information tab for that storage pool do not match, then it indicates that one or more disks are broken and the storage pool is unhealthy.
- Used Allocation
  Displays the number and size of the allocation units used by the aggregates. You can click the aggregate name to view the aggregate details.
- Available Allocation
  Displays the number and size of the allocation units available for the nodes. You can click the node name to view the aggregate details.

Allocated Cache

Displays the size of the allocation units used by the aggregate.

Allocation Units

Displays the number of allocation units used by the aggregate.

Disks

Displays the number of disks contained in the storage pool.

Details

- Storage Pool
  Displays the number of storage pools.
- Total Size
  Displays the total size of the storage pools.

Configuration tab

The Configuration tab displays details about the selected aggregate, such as its cluster node, block type, RAID type, RAID size, and RAID group count:

Overview

- Node
  Displays the name of the node that contains the selected aggregate.
- Block Type
  Displays the block format of the aggregate: either 32-bit or 64-bit.
- RAID Type
  Displays the RAID type (RAID0, RAID4, RAID-DP, RAID-TEC or Mixed RAID).
- RAID Size
  Displays the size of the RAID group.
- RAID Groups
  Displays the number of RAID groups in the aggregate.
- SnapLock Type
  Displays the SnapLock Type of the aggregate.

History area

The History area displays graphs that provide information about the capacity of the selected aggregate.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if the aggregate usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Aggregate Capacity Used (%)

Displays the used capacity in the aggregate and the trend in how aggregate capacity is used based on the usage history as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Capacity Used legend, the Capacity Used graph line is hidden.

Aggregate Capacity Used vs Total Capacity

Displays the trend in how aggregate capacity is used based on the usage history, as well as the used capacity and the total capacity, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Aggregate Capacity Used (%) vs Committed (%)

Displays the trend in how aggregate capacity is used based on the usage history, as well as the committed space as line graphs, as a percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Space Committed legend, the Space Committed graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity

Displays the severity of the event.

Event

Displays the event name.

Triggered Time

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp for when the event was generated is displayed.

Related Devices pane

The Related Devices pane enables you to view the cluster node, volumes, and disks that are related to the aggregate:

Node

Displays the capacity and the health status of the node that contains the aggregate. Capacity indicates the total usable capacity over available capacity.

Aggregates in the Node

Displays the number and capacity of all the aggregates in the cluster node that contains the selected aggregate. The health status of the aggregates is also displayed, based on the highest severity level. For example, if a cluster node contains ten aggregates, five of which display the Warning status and the remaining five of which display the Critical status, then the status displayed is Critical.

Volumes

Displays the number and capacity of the volumes in the selected aggregate. The health status of the volumes is also displayed, based on the highest severity level.

Resource Pool

Displays the resource pools related to the aggregate.

Disks

Displays the number of disks in the selected aggregate.

Related Alerts pane

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

Related tasks

Adding an alert on page 102
Editing aggregate threshold settings on page 128

Related references

Storage Pool dialog box on page 210

Storage Pool dialog box

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

You can view the following storage pool details:

Status

Displays the status of the storage pool, which can be healthy or unhealthy.

Total Allocations

Displays the total allocation units and the size in the storage pool.

Allocation Unit Size

Displays the minimum amount of space in the storage pool that can be allocated to an aggregate.

Disks

Displays the number of disks used to create the storage pool. If the disk count in the storage pool column and the number of disks displayed in the Disk Information tab for that storage pool do not match, then it indicates that one or more disks are broken and the storage pool is unhealthy.

Cache Allocations

- Used Allocations
  Displays the number and size of the allocation units used by the aggregates. You can click the aggregate name to view the aggregate details.
- Available Allocations
  Displays the number and size of the allocation units available for the nodes. You can click the node name to view the aggregate details.

Volumes page

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

- *Command buttons* on page 211
- *Volumes Overview* on page 211
- *Filters pane* on page 215

Command buttons

Edit Thresholds

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

Protect

Displays the following submenus:

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

Restore

Displays the Restore dialog box, which enables you to restore directories or files from one volume at a time. If more than one volume is selected, the **Restore** button is disabled.

Export

Enables you to export the details of all the monitored FlexVol volumes to a comma-separated values (`.csv`) file. However, if you are monitoring Infinite Volumes, you cannot export the details of constituents in an Infinite Volume.

Annotate

Enables you to annotate the selected volume.

Volumes Overview table

The volumes table displays the properties of all the discovered volumes. You can use the column filters to customize the data that is displayed:

Status

Displays the current status of a volume. The status can be Critical (✖), Error (❗), Warning (⚠), or Normal (✓).

You can move the pointer over the status to view more information about the event or events generated for the volume.

If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** link to view more information about the event.

If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

Volume

Displays the name of the volume.

You can move the pointer over a volume to view information such as the qtree quota overcommitted space, status of the last volume move operation, and space allocated in the volume. You can also view the details of related objects such as the SVM to which the volume belongs, the aggregate to which the volume belongs, and all the volumes that belong to this aggregate.

If an SVM with Infinite Volume is monitored, you can view details about the three types of constituents (data, namespace, and namespace mirror) in the SVM with Infinite Volume. The constituent details include the following information:

- Constituent name
- State of the constituent
- Name of the SVM with Infinite Volume to which the constituent belongs
- Junction path of the constituent
- Name of the aggregate that contains the constituent
- Available, used, and total data capacity of the constituent

State

Displays the current state of the volume. The state can be Offline, Online, or Restricted.

Junction Path

Displays the path to which the volume is mounted.

Storage Virtual Machine

Displays the SVM that contains the volume.

Aggregate

Displays the name of the aggregate that contains the volume. You can view more details about the aggregate by clicking the aggregate name.

SnapLock Type

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

In Transition

Indicates whether the volume has completed transition or not. This column is hidden by default.

Protection Role

Displays the protection role of a volume. Protection roles include the following:

- Unprotected

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

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

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

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

You can move your pointer over the protection role for a volume to display a graphical representation of the protection topology for the selected volume, including, if applicable, the source volume, the total number of outgoing SnapMirror relationships, and the total number of outgoing SnapVault relationships. Blue highlighting around the volume indicates the selected volume.

Clicking **View Protection Details** displays the Protection tab of the Volume details page.

Thin Provisioned

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

Available Data Capacity

Displays the amount of space (in KB, MB, GB, and so on) available for data in the aggregate.

Available Data %

Displays the percentage of space available for data in the aggregate.

Used Data Capacity

Displays the amount of space (in KB, MB, GB, and so on) used for data in the volume.

Used Data %

Displays the percentage of space used for data in the volume.

Total Data Capacity

Displays the total data size (in MB, GB, and so on) of the volume.

Storage Class

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

Constituent Role

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

Move Status

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

You can move the pointer over the status to view more information about the volume move operation, such as the source, destination, operation start time, operation end time, current phase of the volume move operation that is in progress, status (in percentage), and estimated end time.

Caching Policy

Displays the caching policy that is associated with the selected volume. The policy provides information about how the flash pool caching occurs for the volume.

Cache policy	Description
Auto	Read caches all the metadata blocks and randomly read user data blocks, and write caches all the randomly overwritten user data blocks.
None	Does not cache any user data or metadata blocks.

Cache policy	Description
All	Read caches all the user data blocks that are read and written. The policy does not perform any write caching.
All-Random Write	This policy is a combination of the All and No Read-Random Write policies and performs the following actions: • Read caches all the user data blocks that are read and written. • Write caches all the randomly overwritten user data blocks.
All Read	Read caches all the metadata, randomly read, and sequentially read user data blocks.
All Read-Random Write	This policy is a combination of the All Read and No Read-Random Write policies and performs the following actions: • Read caches all the metadata, randomly read, and sequentially read user data blocks. • Write caches all the randomly overwritten user data blocks.
All Read Random Write	Read caches all the metadata, randomly read, sequentially read, and randomly written user data blocks.
All Read Random Write-Random Write	This policy is a combination of the All Read Random Write and No Read-Random Write policies and does the following: • Read caches all the metadata, randomly read, and sequentially read, and randomly written user data blocks. • Write caches all the randomly overwritten user data blocks.
Meta	Read caches only metadata blocks.
Meta-Random Write	This policy is a combination of the Meta and No Read-Random Write and does the following: Read caches only
No Read-Random Write	Write caches all the randomly overwritten user data blocks. The policy does not perform any read caching.
Random Read	Read caches all the metadata blocks and randomly read user data blocks.
Random Read-Write	Read caches all the metadata, randomly read, and randomly written user data blocks.

Cache policy	Description
Random Read-Write-Random Write	This policy is a combination of the Random Read Write and No Read-Random Write policies and does the following: • Read caches all the metadata, randomly read, and randomly overwritten user data blocks. • Write caches all the randomly overwritten user data blocks.

Cache Retention Priority

Displays the cache retention priority for the volume. A cache retention priority defines how long the blocks of a volume will be in cache state in a flash pool once they become cold. Cache Retention Priority has the following options:

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

Compression

Indicates whether compression is enabled on the volume. The column displays either Enabled or Disabled. This column is hidden by default.

Deduplication

Displays whether deduplication is enabled on the volume. The column displays either Enabled or Disabled. This column is hidden by default.

Volume Type

Displays which volume type is used. The volume type can be either Read-write or Data-protection. This column is hidden by default.

Cluster

Displays the cluster that contains the destination volume. You can view more details about the cluster by clicking the cluster name. This column is hidden by default.

Cluster Node

Displays the cluster node to which the volume belongs. You can view more details about the cluster node by clicking the node name. This column is hidden by default.

Local Snapshot Policy

Displays the local Snapshot copy policies for the volumes listed. The default policy name is Default. This column is hidden by default.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the volumes list. You can select filters related to the Volume Status, State, and Annotation columns.

Note: The filters specified in the Filters pane override the filters specified for the columns in the volumes list.

Related tasks

Editing volume threshold settings on page 129
Exporting data to CSV files on page 31

Volume details page

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

You must have the OnCommand Administrator or Storage Administrator role.

- *Command buttons* on page 216
- *Capacity tab* on page 217
- *Efficiency tab* on page 220
- *Configuration tab* on page 221
- *Protection tab* on page 222
- *History area* on page 226
- *Events list* on page 227
- *Related Annotations pane* on page 227
- *Related Devices pane* on page 227
- *Related Groups pane* on page 216
- *Related Alerts pane* on page 228

Command buttons

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

- ★: Enables you to add the selected volume to the Favorites dashboard of OnCommand Unified Manager.

Actions

- Add Alert
  Enables you to add an alert to the selected volume.
- Edit Thresholds
  Enables you to modify the threshold settings for the selected volume.
- Annotate
  Enables you to annotate the selected volume.
- Protect
  Enables you to create either SnapMirror or SnapVault relationships for the selected volume.
- Relationship
  Enables you to execute the following protection relationship operations:
    - Edit

    Launches the Edit Relationship dialog box which enables you to change existing SnapMirror policies, schedules, and maximum transfer rates for an existing protection relationship.

  - Abort
    Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer.
  - Quiesce
    Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress must complete before the relationship is quiesced.
  - Break
    Breaks the relationship between the source and destination volumes and changes the destination to a read-write volume.
  - Remove
    Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed. This operation cannot be undone.
  - Resume
    Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used, if one exists.
  - Resynchronize
    Enables you to resynchronize a previously broken relationship.
  - Initialize/Update
    Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform a manual update if the relationship is already initialized.
  - Reverse Resync
    Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by the contents on the destination, and any data that is newer than the data on the common Snapshot copy is deleted.
- Restore
  Enables you to restore data from one volume to another volume.

View Volumes

Enables you to navigate to the Volumes page.

Capacity tab

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

Capacity

Details the display capacity of the volume:

- Snapshot Overflow
  Displays the data space that is consumed by the Snapshot copies.
- Used
  Displays the space used by data in the volume.

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

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

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

- Data graph
  Displays the total data capacity and the used data capacity of the volume.
  If autogrow is enabled, the data graph also displays the space available in the aggregate. The data graph displays the effective storage space that can be used by data in the volume, which can be one of the following:

  - Actual data capacity of the volume for the following conditions:

    - Autogrow is disabled.

    - Autogrow-enabled volume has reached the maximum size.

    - Autogrow-enabled thickly provisioned volume cannot grow further.

  - Data capacity of the volume after considering the maximum volume size (for thinly provisioned volumes and for thickly provisioned volumes when the aggregate has space for the volume to reach maximum size)

  - Data capacity of the volume after considering the next possible autogrow increment (for thickly provisioned volumes that can have at least one autogrow increment)

- Snapshot copies graph
  This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

Autogrow

Displays whether the FlexVol volume automatically grows when it is out of space.

Space Guarantee

Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate. These blocks are then guaranteed to be available for writes to files in the volume. The space guarantee can be set to one of the following:

None

No space guarantee is configured for the volume.

File

Full size of sparsely written files (for example, LUNs) is guaranteed.

Volume

Full size of the volume is guaranteed.

Partial

The FlexCache volume reserves space based on its size. If the FlexCache volume’s size is 100 MB or more, the minimum space guarantee is set to 100 MB by default. If the FlexCache volume’s size is less than 100 MB, the minimum space guarantee is set to the FlexCache volume’s size. If the FlexCache volume’s size is grown later, the minimum space guarantee is not incremented.

Note: The space guarantee is Partial when the volume is of type Data-Cache.

Total Capacity

Displays the total capacity in the volume.

Data Capacity

Displays the amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

Snapshot Reserve

Displays the amount of space used by the Snapshot copies (used capacity) and amount of space available for Snapshot copies (free capacity) in the volume.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the Snapshot copies (used capacity) and the amount of space that is available in the volume but cannot be used for making Snapshot copies (unusable capacity) because of aggregate capacity issues is displayed.

For volumes in a cluster running Data ONTAP 8.1.x, if the Snapshot used reserve is less than 1%, the **Snapshot Reserve Used** field might display a value of 0%, even when there is some used data.

Volume Thresholds

Displays the following volume capacity thresholds:

- Nearly Full Threshold
  Specifies the percentage at which a volume is nearly full.
- Full Threshold
  Specifies the percentage at which a volume is full.

Other Details

- Autogrow Max Size
  Displays the maximum size up to which the FlexVol volume can automatically grow. The default value is 120% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.
- Autogrow Increment Size
  Displays the increment size using which the size of the FlexVol volume increases every time the volume is automatically grown. The default is 5% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.
- Qtree Quota Committed Capacity
  Displays the space reserved in the quotas.
- Qtree Quota Overcommitted Capacity
  Displays the amount of space that can be used before the system generates the Volume Qtree Quota Overcommitted event.
- Fractional Reserve
  Controls the size of the overwrite reserve. By default, the fractional reserve is set to 100, indicating that 100 percent of the required reserved space is reserved so that the objects are fully protected for overwrites. If the fractional reserve is less than 100

percent, the reserved space for all the space-reserved files in that volume is reduced to the fractional reserve percentage.

- Snapshot Daily Growth Rate

  Displays the change (in percentage, or in KB, MB, GB, and so on) that occurs every 24 hours in the Snapshot copies in the selected volume.

- Snapshot Days to Full

  Displays the estimated number of days remaining before the space reserved for the Snapshot copies in the volume reaches the specified threshold.

  The Snapshot Days to Full field displays a Not Applicable value when the growth rate of the Snapshot copies in the volume is zero or negative, or when there is insufficient data to calculate the growth rate.

- Snapshot Autodelete

  Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.

- Snapshot Copies

  Displays information about the Snapshot copies in the volume.

  For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

  For volumes in a cluster running Data ONTAP 8.1.x, the View link is displayed. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

Volume Move

Displays the status of either the current or the last volume move operation that was performed on the volume, and other details, such as the current phase of the volume move operation which is in progress, source aggregate, destination aggregate, start time, end time, and estimated end time.

Also displays the number of volume move operations that are performed on the selected volume. You can view more information about the volume move operations by clicking the **Volume Move History** link.

Efficiency tab

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

Deduplication

- Enabled

  Specifies whether deduplication is enabled or disabled on a volume.

- Space Savings

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

- Last Run

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

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

- Mode

Specifies whether the deduplication operation enabled on a volume is a manual, scheduled, or policy-based operation. If the mode is set to Scheduled, the operation schedule is displayed, and if the mode is set to a policy, the policy name is displayed.

- Status
  Displays the current status of the deduplication operation. The status can be Idle, Initializing, Active, Undoing, Pending, Downgrading, or Disabled.
- Type
  Specifies the type of deduplication operation running on the volume. If the volume is in a SnapVault relationship, the type displayed is SnapVault. For any other volume, the type is displayed as Regular.

Compression

- Enabled
  Specifies whether compression is enabled or disabled on a volume.
- Space Savings
  Displays the amount of space saved (in percentage, or in KB, MB, GB, and so on) in a volume by using compression.

Configuration tab

The Configuration tab displays details about the selected volume, such as the export policy, RAID type, capacity and storage efficiency related features of the volume:

Overview

- Full Name
  Displays the full name of the volume.
- Aggregate
  Displays the name of the aggregate that contains the volume.
- Storage Virtual Machine
  Displays the name of the Storage Virtual Machine (SVM) that contains the volume.
- Junction Path
  Displays the status of the path, which can be active or inactive. The path in the SVM to which the volume is mounted is also displayed. You can click the **History** link to view the most recent five changes to the junction path.
- Export policy
  Displays the name of the export policy that is created for the volume. You can click the link to view details about the export policies, authentication protocols, and access enabled on the volumes that belong to the SVM.
- Type
  Displays the type of the selected volume. The volume type can be Read-write, Load-sharing, Data-Protection, Data-cache, or Temporary.
- Style
  Displays the volume style, which is FlexVol.
- RAID Type
  Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, or RAID-TEC.
- SnapLock Type
  Displays the SnapLock Type of the aggregate that contains the volume.

- SnapLock Expiry
  Displays the expiry date of SnapLock volume.

Capacity

- Thin Provisioning
  Displays whether thin provisioning is configured for the volume.
- Autogrow
  Displays whether the flexible volume grows automatically within an aggregate.
- Snapshot Autodelete
  Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.
- Quotas
  Specifies whether the quotas are enabled for the volume.

Efficiency

- Deduplication
  Specifies whether deduplication is enabled or disabled for the selected volume.
- Compression
  Specifies whether compression is enabled or disabled for the selected volume.

Protection

- Snapshot Copies
  Specifies whether automatic Snapshot copies are enabled or disabled.

Protection tab

The Protection tab displays protection details about the selected volume, such as lag information, relationship type, and topology of the relationship.

Summary

Displays SnapMirror and SnapVault relationships properties for a selected volume. For any other relationship type, only the Relationship Type property is displayed. If a primary volume is selected, only the Managed and Local Snapshot copy Policy are displayed. Properties displayed for SnapMirror and SnapVault relationships include the following:

- Source Volume
  Displays the name of the selected volume's source if the selected volume is a destination.
- Lag Status
  Displays the update or transfer lag status for a protection relationship. The status can be Error, Warning, or Critical.
- Lag Duration
  Displays the time by which the data on the mirror lags behind the source.
- Last Successful Update
  Displays the date and time of the most recent successful protection update.
- Storage Service Member
  Displays either Yes or No to indicate whether or not the volume belongs to and is managed by a storage service.
- Version Flexible Replication

Displays either Yes, Yes with backup option, or None. Yes indicates that SnapMirror replication is possible even if source and destination volumes are running different versions of Data ONTAP 8.3 or later. Yes with backup option indicates the implementation of SnapMirror protection with the ability to retain multiple versions of backup copies on the destination. None indicates that Version Flexible Replication is not enabled.

- Relationship Capability

  Indicates the Data ONTAP capabilities available to the protection relationship. The relationship capability is either prior to Data ONTAP 8.2, or 8.2 and later. A relationship capability of prior to 8.2 means that relationships have not been upgraded to Data ONTAP 8.2 on both the destination and the source clusters, and cannot take advantage of new or improved protection features available in Data ONTAP 8.2 and later. A relationship capability of 8.2 and later means the destination and source clusters are using Data ONTAP 8.2 or later, and can take advantage of improved protection features.

- Protection Service

  Displays the name of the protection service if the relationship is managed by a protection partner application.

- Relationship Type

  Displays any relationship type, including SnapMirror or SnapVault.

- Relationship State

  Displays the state of the SnapMirror or SnapVault relationship. The state can be Uninitialized, SnapMirrored, or Broken-Off. If a source volume is selected, the relationship state is not applicable and is not displayed.

- Transfer Status

  Displays the transfer status for the protection relationship. The transfer status can be one of the following:

  - Idle

    Transfers are enabled and no transfer is in progress.

  - Transferring

    SnapMirror transfers are enabled and a transfer is in progress.

  - Checking

    The destination volume is undergoing a diagnostic check and no transfer is in progress. This applies only to SnapMirror relationships that have the relationship-control-plane field set to v1.

  - Quiescing

    A SnapMirror transfer is in progress. Additional transfers are disabled.

  - Quiesced

    SnapMirror transfers are disabled. No transfer is in progress.

  - Queued

    SnapMirror transfers are enabled. No transfers are in progress.

  - Preparing

    SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.

  - Finalizing

    SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

    - Aborting
      SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

- Max Transfer Rate
  Displays the maximum transfer rate for the relationship. The maximum transfer rate can be can be a numerical value in either kilobytes per second (Kbps), Megabytes per second (Mbps), Gigabytes per second (Gbps), or Terabytes per second (TBps). If No Limit is displayed, the baseline transfer between relationships is unlimited.

- SnapMirror Policy
  Displays the protection policy for the volume. DPDefault indicates the default SnapMirror protection policy, and XDPDefault indicates the default SnapVault policy. You can click the policy name to view details associated with that policy, including the following information:
    - Transfer priority
    - Ignore access time setting
    - Tries limit
    - Comments
    - SnapMirror labels
    - Retention settings
    - Actual Snapshot copies
    - Preserve Snapshot copies
    - Retention warning threshold
    - Snapshot copies with no retention settings

  In a cascading SnapVault relationship where the source is a data protection (DP) volume, only the rule “sm_created” applies.

- Update Schedule
  Displays the SnapMirror schedule assigned to the relationship. Positioning your cursor over the information icon displays the schedule details.

- Local Snapshot Policy
  Displays the Snapshot copy policy for the volume. The policy is Default, None, or any name given to a custom policy.

Views

Displays the protection topology of the selected volume. The topology includes graphical representations of all volumes that are related to the selected volume. The selected volume is indicated by a dark gray border, and lines between volumes in the topology indicate the protection relationship type. Double lines specify a SnapMirror relationship, and a single line specifies a SnapVault relationship. The direction of the relationships in the topology are displayed from left to right, with the source of each relationship on the left and the destination on the right.

Right-clicking a volume displays a menu from which you can choose either to protect the volume or restore data to it.

Right-clicking a relationship displays a menu from which you can choose to either edit, abort, quiesce, break, remove, or resume a relationship. The menus will not display in the following instances:

- If RBAC settings do not allow this action, for example, if you have only operator privileges
- When the volume ID is unknown, for example, when you have a intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 cluster volume

Clicking another volume in the topology selects and displays information for that volume.

A question mark (?) in the upper-left corner of a volume indicates that either the volume is missing or that it has not yet been discovered. It might also indicate that the capacity information is missing. Positioning your cursor over the question mark displays additional information, including suggestions for remedial action.

The topology displays information about volume capacity, lag, Snapshot copies, and last successful data transfer if it conforms to one of several common topology templates. If a topology does not conform to one of those templates, information about volume lag and last successful data transfer is displayed in a relationship table under the topology. In that case, the highlighted row in the table indicates the selected volume, and, in the topology view, bold lines with a blue dot indicate the relationship between the selected volume and its source volume.

Topology views include the following information:

- Capacity
  Displays the total amount of capacity used by the volume. Positioning your cursor over a volume in the topology displays the current warning and critical threshold settings for that volume in the Current Threshold Settings dialog box. You can also edit the threshold settings by clicking the **Edit Thresholds** link in the Current Threshold Settings dialog box. Clearing the **Capacity** check box hides all capacity information for all volumes in the topology.
- Lag
  Displays the lag duration and the lag status of the incoming protection relationships. Clearing the **Lag** check box hides all lag information for all volumes in the topology. When the **Lag** check box is dimmed, then the lag information for the selected volume is displayed in the relationship table below the topology, as well as the lag information for all related volumes.
- Snapshot
  Displays the number of Snapshot copies available for a volume. Clearing the **Snapshot** check box hides all Snapshot copy information for all volumes in the topology.

  Clicking a Snapshot copy icon (❐) displays the Snapshot copy list for a volume. The Snapshot copy count displayed next to the icon is updated every 15 minutes; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon. If you are running Data ONTAP 8.1, the Snapshot copy count is not displayed in the topology.
- Last Successful Transfer
  Displays the amount, duration, time, and date of the last successful data transfer. When the **Last Successful Transfer** check box is dimmed, then the last successful transfer information for the selected volume is displayed in the relationship table below the topology, as well as the last successful transfer information for all related volumes.

History

Displays in a graph the history of incoming SnapMirror and SnapVault protection relationships for the selected volume when you are using Data ONTAP 8.2 or later relationship capabilities. No historical data is collected for relationship capabilities earlier than Data ONTAP 8.2. There are three history graphs available: incoming relationship transfer duration, incoming relationship lag size, and incoming relationship transferred size. History information is displayed only when you select a destination volume. If you select a primary volume, the graphs are empty, and the message `No data found` is displayed.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if large amounts of data are being transferred at the same time of the day or week, or if the lag warning or lag error threshold is consistently being breached, you can take the appropriate action.

History graphs display the following information:

Relationship Transfer Duration

Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum transfer duration reached in the duration period shown in the x axis. You can view the details of specific points on the graph by positioning your cursor over the area of interest.

Relationship Lag Duration

Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum lag duration reached in the duration period shown in the x axis. The horizontal orange line on the graph depicts the lag error threshold, and the horizontal yellow line depicts the lag warning threshold. Positioning your cursor over these lines displays the threshold setting. The horizontal blue line depicts the lag duration. You can view the details for specific points on the graph by positioning your cursor over an area of interest.

Relationship Transferred Size

Displays bytes, kilobytes, megabytes, and so on, on the vertical (y) axis depending on the transfer size, and displays days, months, or years on the horizontal (x) axis depending on the selected time period. The upper value on the y axis indicates the maximum transfer size reached in the duration period shown in the x axis. You can view the details for specific points on the graph by positioning your cursor over an area of interest.

History area

The History area displays graphs that provide information about the capacity and space reservations of the selected volume.

Graphs might be empty and the message `No data found` displayed when the data or the state of the volume remains unchanged for a period of time.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends—for example, if the volume usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Volume Capacity Used

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

Volume Capacity Used vs Total

Displays the trend in how volume capacity is used based on the usage history, as well as the used capacity, total capacity, and details of the space savings from deduplication and compression, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Volume Capacity Used (%)

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

Snapshot Capacity Used (%)

Displays the Snapshot reserve and Snapshot warning threshold as line graphs, and the capacity used by the Snapshot copies as an area graph, in percentage, on the vertical (y) axis. The Snapshot overflow is represented with different colors. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Snapshot Reserve legend, the Snapshot Reserve graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity

Displays the severity of the event.

Event

Displays the event name.

Triggered Time

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp when the event was generated is displayed.

Related Annotations pane

The Related Annotations pane enables you to view annotation details associated with the selected volume. The details include the annotation name and the annotation values that are applied to the volume. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view and navigate to the SVMs, aggregates, qtrees, LUNs, and Snapshot copies that are related to the volume:

Storage Virtual Machine

Displays the capacity and the health status of the SVM that contains the selected volume.

Aggregate

Displays the capacity and the health status of the aggregate that contains the selected volume.

Volumes in the Aggregate

Displays the number and capacity of all the volumes that belong to the parent aggregate of the selected volume. The health status of the volumes is also displayed, based on the highest severity level. For example, if an aggregate contains ten volumes, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.

Qtrees

Displays the number of qtrees that the selected volume contains and the capacity of qtrees with quota that the selected volume contains. The capacity of the qtrees with quota is displayed in relation to the volume data capacity. The health status of the qtrees is also displayed, based on the highest severity level. For example, if a volume has ten qtrees, five with Warning status and the remaining five with Critical status, then the status displayed is Critical.

NFS Exports

Displays the number and status of the NFS exports associated with the volume.

CIFS Shares

Displays the number and status of the CIFS shares.

LUNs

Displays the number and total size of all the LUNs in the selected volume. The health status of the LUNs is also displayed, based on the highest severity level.

User and Group Quotas

Displays the number and status of the user and user group quotas associated with the volume and its qtrees.

FlexClone Volumes

Displays the number and capacity of all the cloned volumes of the selected volume. The number and capacity are displayed only if the selected volume contains any cloned volumes.

Parent Volume

Displays the name and capacity of the parent volume of a selected FlexClone volume. The parent volume is displayed only if the selected volume is a FlexClone volume.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected volume.

Related Alerts pane

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

Related tasks

Adding an alert on page 102
Editing volume threshold settings on page 129

Export Policy Rules dialog box

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

Index

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

Access Protocols

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

Client Match

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

Read Only Access

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

Read Write Access

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

Related references

Storage Virtual Machine details page on page 185

Manage Data Sources page

The Manage Data Sources page enables you to view the statuses of data sources (clusters) that are added to the Unified Manager database. You can also view other details, such as the operation supported by the data source, state of the current operation, start and end time of the operation, and description of the operation.

You must have the OnCommand Administrator or Storage Administrator role.

- *Command buttons* on page 229
- *Data Sources list* on page 230
- *Filters Pane* on page 230

Command buttons

The command buttons enable you to perform the following tasks for a selected data source:

Remove

Enables you to remove a data source if the discovery operation fails.

This button is disabled if the data source is discovered successfully and is polling data.

Rediscover

Enables you to restart the discovery operation in case of discovery failures.

In case of poll failures, you should use the **Rediscover** button from the Cluster details page.

Refresh List

Refreshes the data sources list and the properties associated with the data source.

Data sources table

The Data sources table displays a list of the data sources that are added to the Unified Manager database. You can use column sorting and filtering to customize the data sources that are displayed.

Status

Displays the current discovery status of the data source. The status can be Failed (!), Completed (✓), or In Progress (◌).

You can move the pointer over the status to view more information about the event or events generated for the data source.

If the status of the data source is based on a single event, you can view information such as the event name, time and date when the event was generated, the name of the administrator to whom the event is assigned, and the cause of the event. You can click **View Details** to view more information about the event.

Name

Displays the name of the data source (cluster).

Operation

Displays the current operation that is supported by the data source.

The following operations are supported by the data source:

- Discovery
  Specifies the operation when the data source is being discovered.
- Poll
  Specifies the operation when the data source is successfully discovered and has started sampling data.
- Deletion
  Specifies the operation when the data source (cluster) is deleted from the respective storage objects list (Clusters page).

State

Displays the state of the current operation supported by the data source. The state can be Failed, Completed, or In Progress.

Start Time

Displays the operation start time.

End time

Displays the operation end time.

Description

Displays information about the current operation and operation state of the data source. The description also provides information about when the last failed or successful operation occurred.

For example, if the operation is Poll and the operation state is Failed, the following information is displayed: `Monitoring cycle failed. Last success was 4 hours 30 minutes ago.`

Filters pane

The Filters pane enables you to set filters to customize the display of information in the data sources list.

Related tasks

Restarting cluster discovery on page 146
Removing clusters on page 25

Snapshot Copies on a Volume dialog box

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

List view

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

Snapshot Copy

Displays the name of the Snapshot copy.

Used Space %

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

Total Size

Displays the total size of the Snapshot copy.

Created Time

Displays the timestamp when the Snapshot copy was created.

Dependency

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

Command buttons

The command buttons enable you to perform the following tasks:

Calculate

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

Delete Selected

Deletes one or more Snapshot copies.

Close

Closes the Snapshot Copies on a Volume dialog box.

Recalculate

Enables you to calculate the space that can be reclaimed by deleting the selected Snapshot copies.

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

Related tasks

Deleting Snapshot copies on page 152
Calculating reclaimable space for Snapshot copies on page 153

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 configuration. Discovering a connectivity issue early enables you to manage your MetroCluster configurations effectively.

Parts of a fabric MetroCluster configuration

As you plan your MetroCluster configuration, you should understand the hardware components and how they interconnect.

DR groups

A fabric Metrocluster configuration consists of one or two DR groups, depending on the number of nodes in the MetroCluster configuration. Each DR group consists of four nodes.

- An eight-node MetroCluster consists of two DR groups.
- A four-node MetroCluster consists of one DR group.

The following illustration shows the organization of the nodes an eight-node MetroCluster configuration:

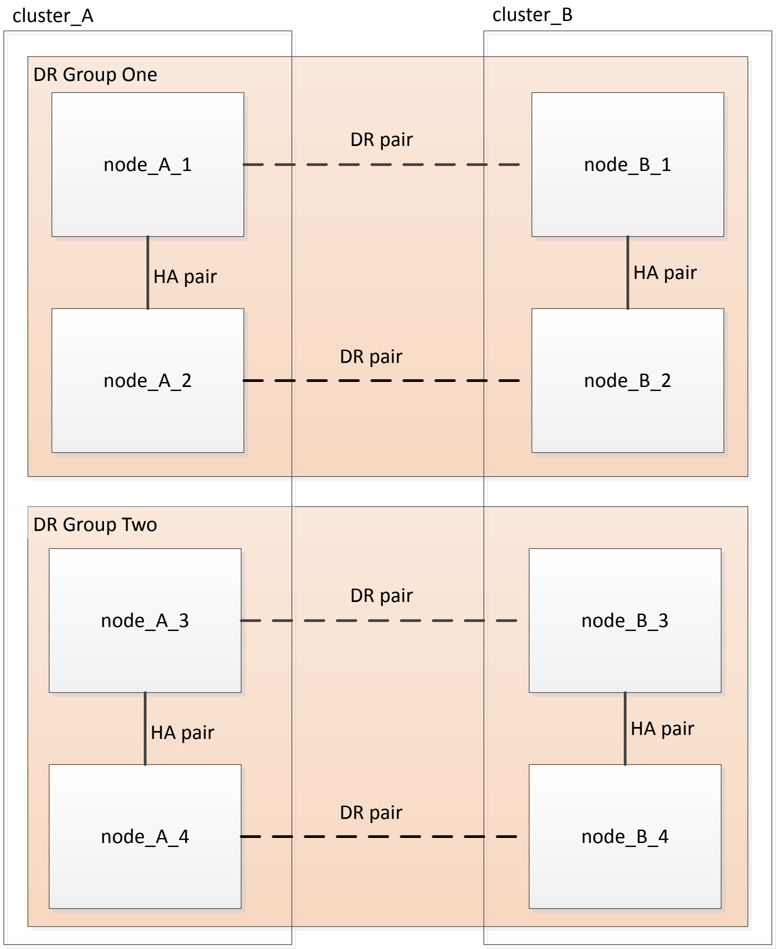

The following illustration shows the organization of the nodes an four-node MetroCluster configuration:

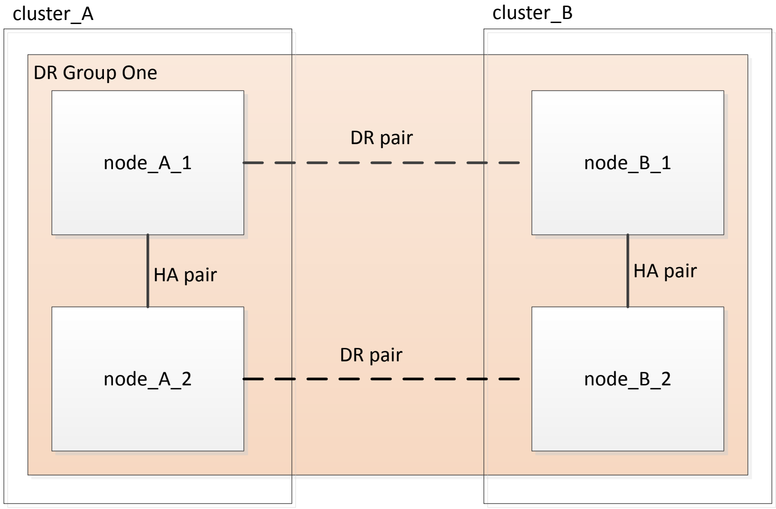

Key hardware elements

The MetroCluster configuration includes the following key hardware elements:

- Storage controllers

  The storage controllers are not connected directly to the storage but connect to two redundant FC switch fabrics.

  Each storage controller is configured as a DR partner to a storage controller on the partner site. When the MetroCluster is enabled, the system will automatically pair the two nodes with lowest system IDs in each of the two clusters as DR partners. In a four-node MetroCluster, there are two DR partner pairs. The second DR pair is created from the two nodes with higher system IDs. Likewise, in an eight-node MetroCluster, the third and fourth DR pairs are created from the next higher system IDs.
- FC-to-SAS bridges

  The FC-to-SAS bridges connect the SAS storage stacks to the FC switches, providing bridging between the two protocols.
- FC switches

  The FC switches provide the long-haul backbone ISL between the two sites. The switches provide the two storage fabrics that allow data mirroring to the remote storage pools.
- Cluster peering network

  The cluster peering network provides connectivity for mirroring of the Storage Virtual Machine (SVM) configuration. The configuration of all SVMs on one cluster is mirrored to the partner cluster.

Eight-node fabric MetroCluster configuration

- The configuration consists of two clusters, one at each geographically separated site.
- cluster_A is located at one MetroCluster site.
- cluster_B is located at the second MetroCluster site.
- Each site has one stack of SAS storage.

  Additional storage stacks are supported, but only one is shown at each site.

- The HA pairs are configured as switchless clusters, without cluster interconnect switches.
  A switched configuration is supported but not shown.

The configuration includes the following connections:

- FC connections from each controller's HBAs and FC-VI adapters to each of the FC switches
- An FC connection from each FC-to-SAS bridge to an FC switch
- SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge
- An HA interconnect between each controller in the local HA pair
  If the controllers support a single-chassis HA pair, the HA interconnect is internal, occurring through the backplane, meaning an external interconnect is not required.
- Ethernet connections from the controllers to the customer-provided network used for cluster peering
  SVM configuration is replicated over the cluster peering network.
- A cluster interconnect between each controller in the local HA pair
  If the controllers are configured as a switched cluster, each controller would connect to two cluster interconnect switches.

Four-node fabric MetroCluster configuration

The following illustration shows a simplified view of a four-node fabric MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.

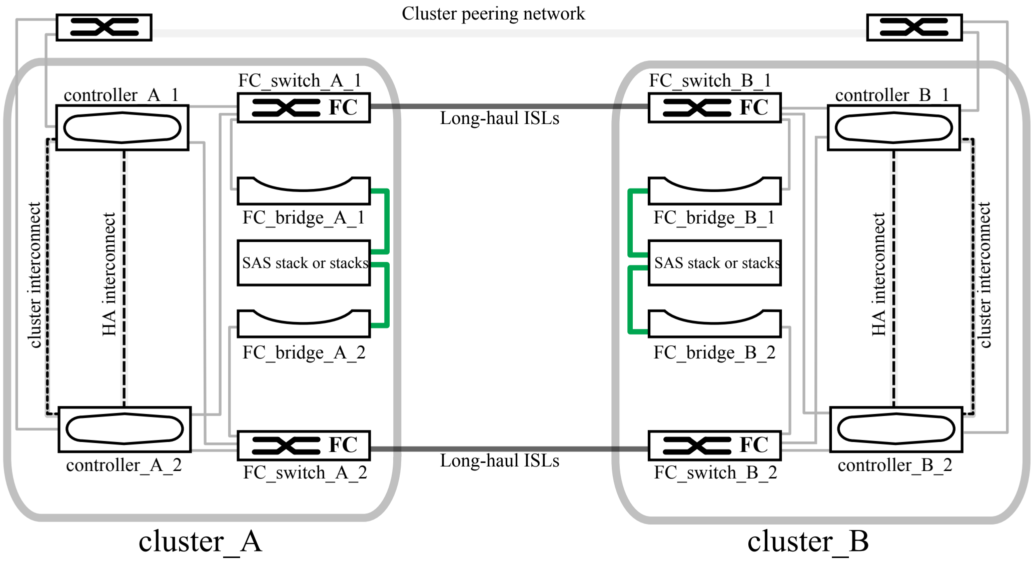

The following illustration shows a more detailed view of the connectivity in a single MetroCluster cluster (both clusters have the same configuration):

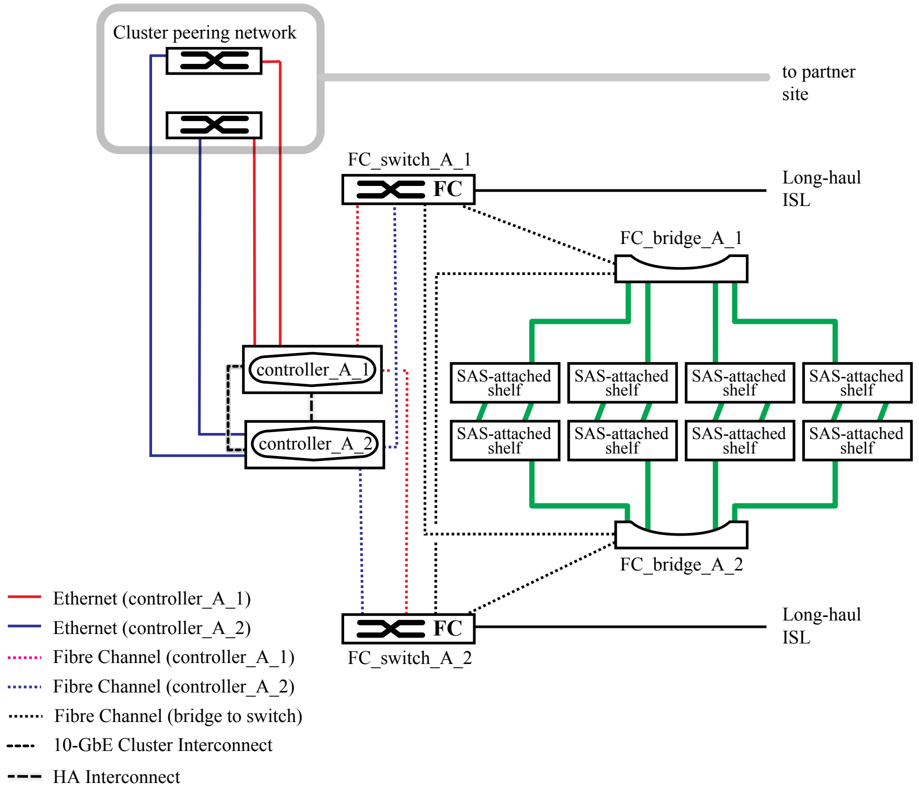

Two-node fabric MetroCluster configuration

The following illustration shows a simplified view of a two-node fabric MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.

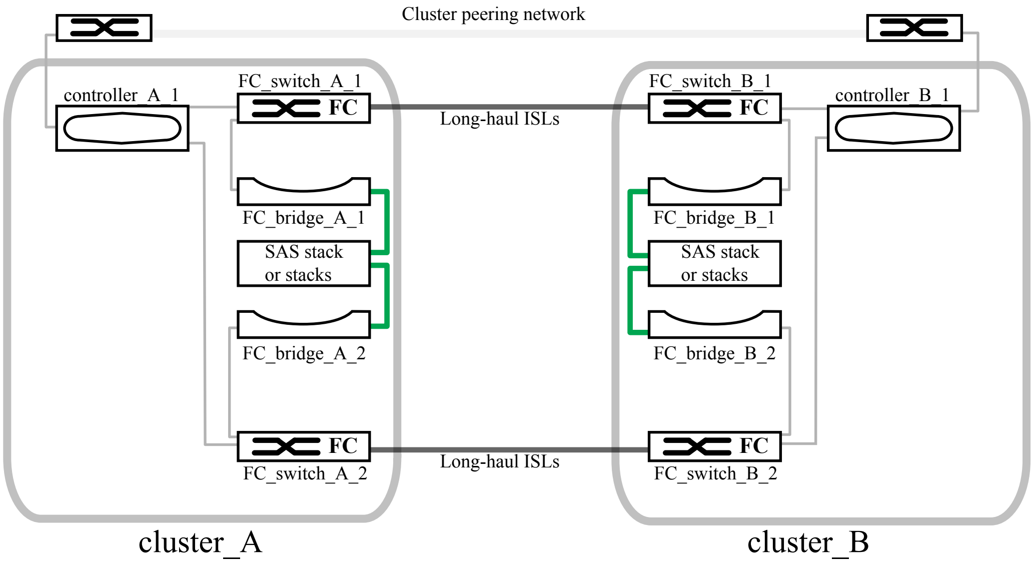

- The configuration consists of two clusters, one at each geographically separated site.
- cluster_A is located at one MetroCluster site.

- cluster_B is located at the second MetroCluster site.
- Each site has one stack of SAS storage.
  Additional storage stacks are supported, but only one is shown at each site.

Note: In the two-node configuration, the nodes are not configured as an HA pair.

The following illustration shows a more detailed view of the connectivity in a single MetroCluster cluster (both clusters have the same configuration):

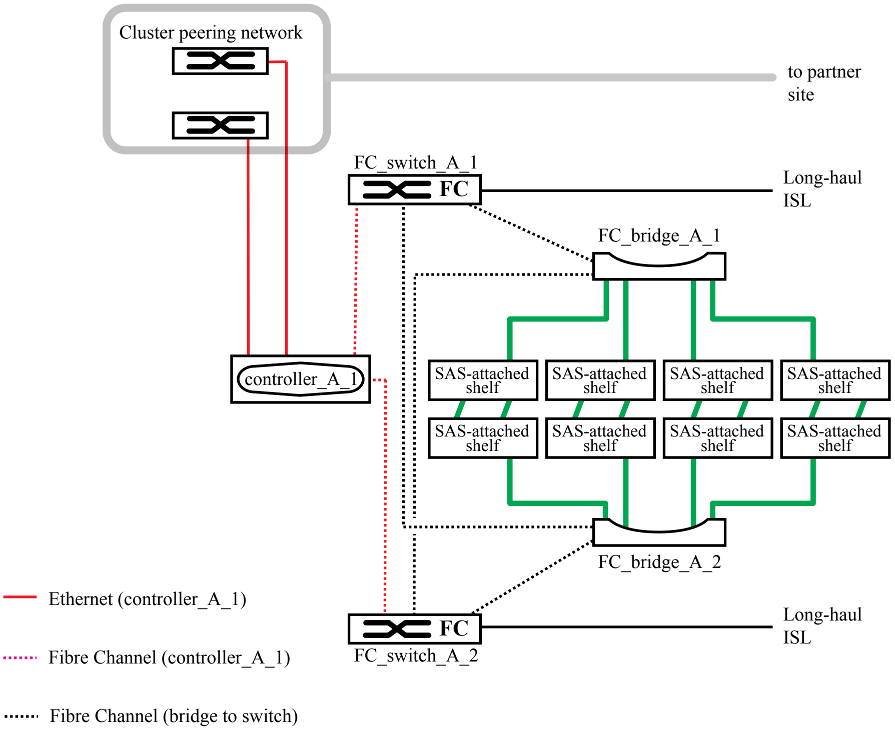

The configuration includes the following connections:

- FC connections between the FC-VI adapter on each controller module
- FC connections from each controller module's HBAs to FC-to-SAS bridge for each SAS shelf stack
- SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge
- Ethernet connections from the controllers to the customer-provided network used for cluster peering
  SVM configuration is replicated over the cluster peering network.

Parts of a two-node direct-attached MetroCluster configuration

The two-node MetroCluster direct-attached configuration requires a number of parts, including two single-node clusters in which the controller modules afre directly connected to the storage using SAS cables.

The MetroCluster configuration includes the following key hardware elements:

- Storage controllers
  The storage controllers connect directly to the storage using SAS cables.

Each storage controller is configured as a DR partner to a storage controller on the partner site. When the MetroCluster is enabled, the system will automatically pair the two nodes with lowest system IDs in each of the two clusters as DR partners.

  - Copper SAS cables can be used for shorter distances.
  - Optical SAS cables can be used for longer distances.

NetApp Interoperability Matrix Tool
After opening the Interoperability Matrix, you can use the Storage Solution field to select your MetroCluster solution (fabric MetroCluster or stretch MetroCluster, with or without FlexArray). You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- Cluster peering network
  The cluster peering network provides connectivity for mirroring of the Storage Virtual Machine (SVM) configuration. The configuration of all SVMs on one cluster is mirrored to the partner cluster.

Parts of a two-node MetroCluster configuration with FC-to-SAS bridges

As you plan your MetroCluster configuration you should understand the parts of the configuration and how they work together.

The MetroCluster configuration includes the following key hardware elements:

- Storage controllers
  The storage controllers are not connected directly to the storage but connect to FC-to-SAS bridges. The storage controllers are connected to each other by FC cables between each controller's FC-VI adapters.
  Each storage controller is configured as a DR partner to a storage controller on the partner site. When the MetroCluster is enabled, the system will automatically pair the two nodes with lowest system IDs in each of the two clusters as DR partners.
- FC-to-SAS bridges
  The FC-to-SAS bridges connect the SAS storage stacks to the FC switches, providing bridging between the two protocols.
- Cluster peering network
  The cluster peering network provides connectivity for mirroring of the Storage Virtual Machine (SVM) configuration. The configuration of all SVMs on one cluster is mirrored to the partner cluster.

The following illustration shows a simplified view of the MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.

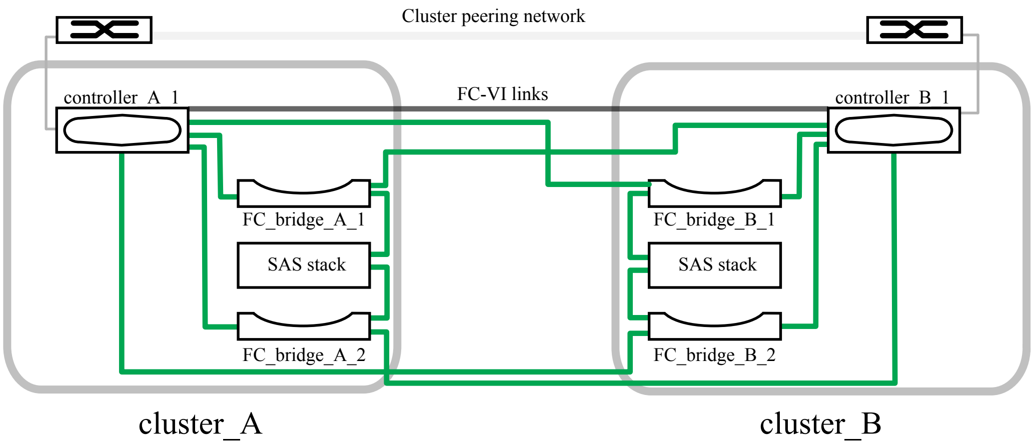

- The configuration consists of two single-node clusters.
- Each site has one stack of SAS storage.

  Note: SAS shelves in MetroCluster configurations are not supported with ACP cabling.

  Additional storage stacks are supported, but only one is shown at each site.

Cluster connectivity status definitions

Connectivity between the clusters in a MetroCluster 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.	● ● ⊗ ● ●

Connectivity status	Description	Icon displayed
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.	Switchover with errors: ● ● ❌ ● ● Switchover successful: ● ● ℹ ● ●

Data mirroring status definitions

MetroCluster 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 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.	🟢 🟢 🟢 🟢 🟢
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.	Switchover with errors: ● ● ❌ ● ● Switchover successful: ● ● ℹ ● ●

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.

Before you begin

- Both the local and remote clusters in the MetroCluster configuration must be added to OnCommand Unified Manager.
- You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

You can use the information displayed in the Cluster 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.

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

The MetroCluster Connectivity tab is displayed only for clusters in a MetroCluster configuration.

Steps

1. Click **Storage > Clusters**.

   A list of all of the monitored clusters is displayed.

2. From the **Clusters** page, click the name of the cluster for which you want to view MetroCluster configuration details.

3. In the **Cluster details** page, click the **MetroCluster Connectivity** tab.

   The topology of the MetroCluster configuration is displayed in the corresponding cluster object area.

After you finish

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

Related references

Cluster details page on page 170
Clusters page on page 168
Cluster connectivity status definitions on page 239

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.

Before you begin

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

About this task

You can use the information displayed in the Cluster details page to rectify any replication issues.

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

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

Steps

1. Click **Storage** > **Clusters.**

   A list of the monitored clusters is displayed.

2. From the **Clusters** page, click the name of the cluster for which you want to view MetroCluster replication details.
3. In the **Clusters** 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.

After you finish

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

Related references

Data mirroring status definitions on page 240
Cluster details page on page 170
Clusters page on page 168

Managing annotations for storage objects

You can create annotations in Unified Manager to annotate storage objects. Annotations enable you to easily identify critical resources, and to take appropriate actions, such as adding critical resources to a group and assigning a group action, or creating a report of annotated resources.

What annotations are

Annotations are dynamically associated with storage objects using annotation rules. When you associate storage objects with predefined annotations, you can filter and view the events that are related to them. You can apply annotations to clusters, volumes, and Storage Virtual Machines (SVMs).

An annotation consists of a name-value pair. Each annotation name can have multiple values and one name-value pair is associated with a storage object through rules. For example, you can create an annotation named “data-center” with the values “Boston” and “Canada”. You can then apply the annotation “data-center” with the value “Boston” to volume v1. When an alert is generated for any event on a volume v1 that is annotated with “data-center”, the generated email indicates the location of the volume, “Boston”, and this enables you to prioritize and resolve the issue.

Related tasks

Viewing the annotation list and details on page 247
Deleting annotations on page 247

Related references

Add Annotation dialog box on page 256
Add Annotation Rule dialog box on page 257

How annotation rules work in Unified Manager

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

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

Conditions

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

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

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

Storage object type	Applicable operands
Volume	• Object name • Owning cluster name • Owning SVM name • Annotations
SVM	• Object name • Owning cluster name • Annotations
Cluster	• Object name • Annotations

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

Example of an annotation rule with conditions

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

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

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

Condition groups

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

Example of an annotation rule with condition groups

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

- Condition group 1
    - Name contains “vol”
    - SVM name is “data_svm”

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

- Condition group 2
    - Name contains “vol”
    - The annotation value of data-priority is “critical”

  This condition group annotates all volumes that include “vol” in their names and that are annotated with the data-priority annotation value as “critical”.

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

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

Description of predefined annotation values

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

Mission critical

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

High

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

Low

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

Adding annotations

You can create custom annotations, and dynamically associate clusters, Storage Virtual Machines (SVMs), and volumes with the annotations by using rules. These rules automatically assign the annotations to storage objects.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Health** > **Administration** > **Manage Annotations**.
2. In the **Annotations** tab, click **Add Annotation**.
3. In the **Add Annotation** dialog box, type a name and description for the annotation.

   You can also add values to annotations while creating annotations.
4. Click **Save and Close**.

Related tasks

Adding values to annotations on page 246

Related references

Manage Annotations page on page 253
Add Annotation dialog box on page 256

Adding values to annotations

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot add values to predefined annotations.

Steps

1. Click **Health** > **Administration** > **Manage Annotations**.
2. From the annotations list in the**Annotations** tab, select the annotation to which you want to add a value.
3. In the **Values** section of the **Annotations** tab, click **Add**.
4. In the **Add Annotation Value** dialog box, specify a value for the annotation.

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

Related tasks

Adding annotations on page 246

Related references

Manage Annotations page on page 253

Deleting annotations

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

Before you begin

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

Steps

1. Click **Health** > **Administration** > **Manage Annotations**.
2. In the **Annotations** tab, select the annotation that you want to delete.

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

Result

The selected annotation and its value is deleted.

Related tasks

Deleting values from annotations on page 248

Related references

Manage Annotations page on page 253

Viewing the annotation list and details

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

Steps

1. Click **Health** > **Administration** > **Manage Annotations**.
2. In the **Annotations** tab, click the annotation name to view the associated details.

Related references

Manage Annotations page on page 253

Deleting values from annotations

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

Before you begin

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

About this task

You cannot delete values from predefined annotations.

Steps

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

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

Related tasks

Deleting annotations on page 247

Related references

Manage Annotations page on page 253

Creating annotation rules

You can create annotation rules to dynamically annotate storage objects such as volumes, clusters, or Storage Virtual Machines (SVMs).

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. Click **Health** > **Administration** > **Manage Annotations**.
2. In the **Annotation Rules** tab, click **Add**.
3. In the **Add Annotation Rules** dialog box, specify a name for the annotation rule.

4. In the **Target Object Type** field, select the type of storage object that you want to annotate.
5. In the **Apply Annotation** field, select the required annotation and annotation value for the annotation.
6. In the **Conditions** section, perform the appropriate action to create a condition, a condition group, or both:

To create...	Do this...
A condition	**a.** Select an operand from the list of operands. **b.** Select either **Contains** or **Is** as the operator. **c.** Enter a value, or select a value from the available list.
A condition group	**a.** Click **Add Condition Group**. **b.** Select an operand from the list of operands. **c.** Select either **Contains** or **Is** as the operator. **d.** Enter a value, or select a value from the available list. **e.** Click **Add condition** to create more conditions if required, and repeat steps a through d for each condition.

7. Click **Add**.

Example of creating an annotation rule

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

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

Related concepts

How annotation rules work in Unified Manager on page 243

Related tasks

Editing annotation rules on page 250
Reordering annotation rules on page 252
Deleting annotation rules on page 252

Related references

Reorder Annotation Rule dialog box on page 258

Adding annotations to storage objects

Unified Manager enables you to manually annotate volumes, clusters, and SVMs without using annotation rules. You can annotate a single storage object or multiple storage objects and specify the required combination of value pair for the annotation.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Add an annotation to required storage objects:

To add annotation to…	Do this…
Clusters	**a.** Click **Storage** > **Clusters**. **b.** Select the required cluster or clusters.
Volumes	**a.** Click **Storage** > **Volumes** > . **b.** Select the required volume or volumes.
SVMs	**a.** Click **Storage** > **Volumes** > **SVM**. **b.** Select the required SVM or SVMs.

2. Click **Annotate**.

Editing annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. Click **Health** > **Administration** > **Manage Annotations**.
2. In the **Annotation Rules** tab, select the annotation rule you want to edit, and then click **Edit**.

3. In the **Edit Annotation Rule** dialog box, change the rule name, annotation name and value, condition groups, and conditions as required.

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

4. Click **Save**.

Related concepts

How annotation rules work in Unified Manager on page 243

Related tasks

Creating annotation rules on page 248
Deleting annotation rules on page 252

Configuring conditions for annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Health** > **Administration** > **Manage Annotations**.
2. In the **Annotation Rules** tab, click **Add**.
3. In the **Add Annotation Rule** dialog box, select one operand from the list of operands.
4. Select an operator for the condition.
5. Enter a value or select one from the available list.
6. Click **Add**.

Example of configuring a condition for an annotation rule

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

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

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

Deleting annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

Related concepts

How annotation rules work in Unified Manager on page 243

Related tasks

Creating annotation rules on page 248
Editing annotation rules on page 250

Reordering annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. Click **Health** > **Administration** > **Manage Annotations**.
2. In the **Annotation Rules** tab, click **Reorder**.
3. In the **Reorder Annotation Rule** dialog box, drag and drop single or multiple rows to rearrange the sequence of the annotation rules.

4. Click **Save**.

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

Description of Annotations windows and dialog boxes

You can view and manage all your annotations from the Manage Annotations page. You can also configure annotation rules for your storage objects from the Annotation Rules tab.

Manage Annotations page

The Manage Annotations page enables you to create annotations in Unified Manager that can be used to annotate storage objects. You can either manually annotate storage objects with an annotation=value pair or configure annotation rules. Storage objects are annotated dynamically based on the annotation you apply.

When you log in as an operator, you will have only read access to the page. You can access the add, edit, or delete buttons in each tab when you log in as a storage administrator or Unified Manager administrator.

- *Annotations tab* on page 253
- *Annotation Rules tab* on page 255

Annotations tab

The Annotations tab enables you to view, create, or delete annotations in Unified Manager.

Annotations list

Displays the names of the predefined and custom annotations. The count of the annotation values associated with each annotation is also displayed. You can click the annotation name to view the details of the annotation.

Summary area

You can view the following details of the selected annotation:

Description

Displays the description provided for the annotation.

Data Type

Displays whether the selected annotation is rule based or a manual annotation.

Created by

Displays the name of the user who created the annotation.

Creation date

Displays the date when the annotation was created.

Annotation=Values Pairs

Displays the list of annotation-value pairs and associated storage objects that are available for the selected annotation.

Value

Displays the name of the annotation=value pair.

Applicable Clusters

Displays the number of clusters that are annotated with a particular annotation=value pair. You can click the number to view the clusters page, which displays a filtered list of the clusters associated with a specific value.

Applicable SVMs

Displays the number of SVMs that are annotated with a particular annotation=value pair.

You can click the number to view the SVMs page, which displays a filtered list of SVMs associated with a specific value.

Applicable Volumes

Displays the number of volumes that are annotated with a particular annotation=value pair. You can click the number to view the volumes page, which displays a filtered list of the volumes associated with a specific value.

Object Associations via Rules

Displays the list of annotation rules and the associated storage objects for the selected annotation.

Rank

Displays the order of the annotation rules to be applied on the storage objects.

Rules

Displays the name of the annotation rule.

Target Object Type

Displays the type of storage object to which the annotation rule is applied.

Annotation=Value Pair

Displays the annotation=value pair applied to the storage object.

Applicable Objects

Displays the count of the storage objects that are annotated based on the annotation rule.

Manual Object Associations

Displays the list of annotations that you have manually configured and associated with storage objects.

Annotation=Value Pair

Displays the name of the manual annotation and the value.

Applicable Clusters

Displays the number of clusters that are annotated with a particular manual annotation value. You can click the number to view the clusters page, which displays a filtered list of the clusters associated with a specific value.

Applicable SVMs

Displays the number of SVMs that are annotated with a particular manual annotation value.

You can click the number to view the SVMs page, which displays a filtered list of SVMs associated with a specific value.

Applicable Volumes

Displays the number of volumes that are annotated with a particular manual annotation value. You can click the number to view the volumes page, which displays a filtered list of the volumes associated with a specific value.

Command buttons

You must have the OnCommand Administrator or Storage Administrator role. For predefined annotations, you cannot add or delete values.

Add Annotation

Opens the Add Annotation dialog box, which enables you to create new custom annotations and assign values to the annotation.

Actions

Enables you to edit the selected annotation description and delete the selected custom annotation. You can delete the annotation only when none of its values are associated with any annotation rules or group rules.

Add

Opens the Add Annotation Value dialog box, which enables you to add a new value for the annotation.

Delete

Enables you to delete the annotation value. You can delete the value only when it is not associated with any annotation rules or group rules.

Edit Rule

Enables you to modify the annotation rules that are configured for the selected annotation.

Annotation Rules tab

The Annotations Rules tab displays the annotation rules you created to annotate storage objects. You can perform tasks such as adding, editing, deleting, or reordering an annotation rule. You can also view the number of storage objects that satisfy the annotation rule.

Command buttons

You must have the OnCommand Administrator or Storage Administrator role.

Add

Displays the Add Annotation Rule dialog box, which enables you to create annotation rules for storage objects.

Edit

Displays the Edit Annotation Rule dialog box, which enables you to reconfigure previously configured annotation rules.

Delete

Deletes the selected annotation rules.

Reorder

Displays the Reorder Annotation Rule dialog box, which enables you to rearrange the order of the annotation rules.

List View

The list view displays, in tabular format, the annotation rules you created in the Unified Manager server. You can use the column filters to customize the data that is displayed. The list view of the Annotation Rules tab and the list view of the Associated Rules section in the Annotation tab contains the following columns:

- Rank
- Name

- Target Object type
- Associated Annotation Value
- Applicable Objects

An additional column is displayed for the Annotation Rules tab, Associated Annotation, which displays the name of the annotation applied to the storage object.

Related tasks

Adding annotations on page 246
Deleting annotations on page 247
Adding values to annotations on page 246
Deleting values from annotations on page 248

Related references

Description of predefined annotation values on page 245

Add Annotation dialog box

The Add Annotation dialog box enables you to create custom annotations that you can associate with clusters, volumes, and Storage Virtual Machines (SVMs) through annotation rules.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation Name

Specifies the name of the annotation. You must enter a unique name for the annotation.

Description

Specifies a meaningful description of the annotation.

Annotation Values

Add

Adds a new value to the selected annotation.

Delete

Deletes the selected value for an annotation.

Command buttons

Save and Close

Saves the new annotation and closes the Add Annotation dialog box dialog box.

Cancel

Closes the Add Annotation dialog box without saving your changes.

Related tasks

Adding annotations on page 246

Related references

Manage Annotations page on page 253

Edit Annotation dialog box

The Edit Annotation dialog box enables you to change the description of an existing annotation.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation Name

Displays the name of the annotation. This field cannot be edited.

Description

Provides a meaningful description of the annotation. You can edit this field when you want to change the current description of the annotation.

Command buttons

Save and Close

Saves the annotation description changes and closes the dialog box.

Cancel

Closes the Edit Annotation dialog box without saving your changes.

Related references

Manage Annotations page on page 253

Add Annotation Rule dialog box

The Add Annotation Rule dialog box enables you to create annotation rules in Unified Manager to dynamically annotate storage objects.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Specifies the name of the annotation rule.

Target Object Type

Specifies the type of storage objects (Storage Virtual Machines (SVMs), volumes, or clusters) that you want to annotate.

Apply Annotation

Specifies the annotation and the value you can use to annotate storage objects when all conditions are met.

Conditions

Specifies conditions that determine which storage objects you can annotate.

Command buttons

Save and Add

Adds the annotation rule you created and enables you to add another annotation rule without closing the dialog box.

Add

Adds the annotation rule and closes the Add Annotation Rule dialog box.

Cancel

Cancels the changes and closes the Add Annotation Rule dialog box.

Add Condition

Adds a condition to define the annotation rule.

Add Condition Group

Adds a condition group to define conditions for the annotation rule.

Related concepts

How annotation rules work in Unified Manager on page 243

Related references

Edit Annotation Rule dialog box on page 258

Edit Annotation Rule dialog box

You can edit the annotation rules you created to add or remove annotations on storage objects.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Displays the name of the annotation rule.

Target Object Type

Displays the type of storage object that you want to annotate. You cannot change the object type.

Apply Annotation

Displays the annotation and the value you can use to annotate storage objects when all conditions are met.

Conditions

Displays the list of conditions for the annotation rule. You can edit the conditions to add or remove the annotation on storage objects.

Command buttons

Save

Saves the changes you made and closes the Edit Annotation Rule dialog box.

Cancel

Closes the Edit Annotation Rule dialog box without saving your changes.

Related concepts

How annotation rules work in Unified Manager on page 243

Related references

Add Annotation Rule dialog box on page 257

Reorder Annotation Rule dialog box

You can use the Reorder Annotation Rule dialog box to specify the order in which you want annotation rules to be applied to storage objects.

Command buttons

You must have the OnCommand Administrator or Storage Administrator role.

Save

Saves the changes you made to the annotation rules and closes the Reorder Annotation Rule dialog box.

Cancel

Closes the Reorder Annotation Rule dialog box without saving the changes you made.

List View

Rank

Displays the order in which the annotation rules will be applied to the storage objects.

Name

Displays the name of the annotation rule.

Target Object Type

Displays the type of storage object to which the annotation rule is applied.

Associated Annotation

Displays the name of the annotation that is applied to the storage object.

Associated Annotation Value

Displays the annotation value for the storage object.

Annotate Cluster dialog box

The Annotate Cluster dialog box enables you to manually annotate storage objects. You can select either a single cluster or multiple clusters and annotate with a specific value pair from the existing list of annotations.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation=Value Pairs

Enables you to select the required annotation for the selected cluster.

Apply

Applies the selected annotation to the cluster.

Cancel

Closes the Annotate Cluster dialog box without saving your changes.

Annotate SVM dialog box

The Annotate SVM dialog box enables you to manually annotate storage objects. You can select either a single SVM or multiple SVMs and annotate with a specific value pair from the existing list of annotations.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation=Value Pairs

Enables you to select the required annotation for the selected SVM.

Apply

Applies the selected annotation to the SVM.

Cancel

Closes the Annotate SVM dialog box without saving your changes.

Annotate Volume dialog box

The Annotate Volume dialog box enables you to manually annotate storage objects. You can select either a single volume or multiple volumes and annotate with a specific value pair from the existing list of annotations.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation=Value Pairs

Enables you to select the required annotation for the selected volume.

Apply

Applies the selected annotation to the volume.

Cancel

Closes the Annotate Volume dialog box without saving your changes.

Managing and monitoring groups

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

Understanding groups

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

What a group is

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

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

Related references

Add Group dialog box on page 272
Add Group Rule dialog box on page 274
Add Group Action dialog box on page 275

How group rules work for groups

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

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

Conditions

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

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

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

Storage object type	Applicable operands
Volume	• Object name • Owning cluster name • Owning SVM name • Annotations
SVM	• Object name • Owning cluster name • Annotations
Cluster	• Object name • Annotations

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

Example of a group rule with conditions

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

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

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

Condition groups

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

Example of a group rule with condition groups

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

- Condition group 1
    - Name contains “vol”
    - SVM name is “data_svm”

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

  This condition group selects all volumes that include “vol” in their names and that are annotated with the data-priority annotation value as “critical”.

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

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

How group actions work on storage objects

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

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

Example of unique group actions

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

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

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

Managing groups of storage objects

You can manage storage objects in your environment by creating groups of storage objects. These storage objects must satisfy the group rules associated with the group.

Adding groups

You can create groups to combine clusters, volumes, and Storage Virtual Machines (SVMs) for ease of management.

Before you begin

SVMsYou must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. Click **Health** > **Administration** > **Manage Groups**.
2. In the **Groups** tab, click **Add**.
3. In the **Add Group** dialog box, enter a name and description for the group.

   The group name must be unique.
4. Click **Add**.

Related tasks

Adding group rules on page 264
Adding group actions on page 267

Deleting groups

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

Before you begin

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

Steps

1. Click **Health** > **Administrator** > **Manage Groups**.
2. In the **Groups** tab, select the group that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box , confirm the deletion by clicking **Yes**.

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

Editing groups

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

Adding group rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

To create....	Do this...
A condition	**a.** Select an operand from the list of operands. **b.** Select either **Contains** or **Is** as the operator. **c.** Enter a value, or select a value from the available list.

To create....	Do this...
A condition group	**a.** Click **Add Condition Group** **b.** Select an operand from the list of operands. **c.** Select either **Contains** or **Is** as the operator. **d.** Enter a value, or select a value from the available list. **e.** Click **Add condition** to create more conditions if required, and repeat steps a through d for each condition.

7. Click **Add**.

Example for creating a group rule

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

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

Related concepts

How group rules work for groups on page 260

Related tasks

Editing group rules on page 266
Deleting group rules on page 266

Editing group rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

   Note: You cannot change the target object type for a group rule.
4. Click **Save**.

Related concepts

How group rules work for groups on page 260

Related tasks

Adding group rules on page 264
Deleting group rules on page 266

Deleting group rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. Click **Health** > **Administration** > **Manage Groups**.
2. In the **Group Rules** tab, select the group rule that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Related concepts

How group rules work for groups on page 260

Related tasks

Adding group rules on page 264
Editing group rules on page 266

Configuring conditions for group rules

You can configure one or more conditions to create group rules in Unified Manager that are applied on the storage objects. The storage objects that satisfy the group rule are combined into a group.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Health > Administration > Manage Groups > Group Rules** tab.
2. Click **Add**.
3. In the **Add Group Rule** dialog box, select one operand from the list of operands.
4. Select an operator for the condition.
5. Enter a required value or select one from the available list.
6. Click **Add**.

Example of configuring a condition for a group rule

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

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

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

Adding group actions

You can configure group actions that you want to apply to storage objects in a group in OnCommand Unified Manager by using the Manage Groups page. Configuring actions for a group enables you to save time, because you do not have to add these actions to each object individually.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Health > Administration > Manage Groups**.
2. In the **Group Actions** tab, click **Add**.
3. In the **Add Group Action** dialog box, enter a name and description for the action.
4. From the **Group** menu, select a group for which you want to configure the action.

5. From the **Action Type** menu , select an action type.

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

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

Related tasks

Editing group actions on page 268

Editing group actions

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

Related tasks

Deleting group actions on page 269

Configuring volume thresholds for groups

You can configure group-level volume thresholds by using Unified Manager. You can configure volume thresholds for capacity, Snapshot copies, qtree quotas, growth, and inodes.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

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

Deleting group actions

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. Click **Health > Administration > Manage Groups**.
2. In the **Group Actions** tab, select the group action that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Reordering group actions

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

Description of groups windows and dialog boxes

You can use the Manage Groups page to view and manage all your groups. You can also configure group rules and actions for your storage objects from the Group Rules tab and Group Actions tab.

Manage Groups page

The Manage Groups page enables you to create groups in Unified Manager to easily manage storage objects. A group is a dynamic collection of storage objects (clusters, volumes, and SVMs), which is defined by the group rules you create for the group.

The Manage Groups page includes tabs that enable you to add, delete, or edit a group, group rules, and actions. When you log in as an operator, you will have only read access to the page. You can access the add, edit, or delete buttons in each tab when you log in as a storage administrator or Unified Manager administrator.

- *Groups tab* on page 270
- *Group Rules tab* on page 271
- *Group Actions tab* on page 271

Groups tab

The Groups tab displays the name and description of the groups you created. You can perform tasks such as adding, editing, or deleting a group. The tab also displays the number of group rules and group actions associated with a group, the number of clusters, SVMs, and volumes in the group.

Command buttons

Add

Displays the Add Group dialog box, which enables you to add a group and provide a name and description for the group.

You can also apply group rules later to the group to include storage objects.

Edit

Displays the Edit Group dialog box, which enables you to edit the name and description for the selected group.

Delete

Deletes the selected group.

List view

The list view displays, in tabular format, the groups you created in Unified Manager. You can use the column filters to customize the data that is displayed. By default, the list is sorted by group name.

Name

Displays the name of the group.

Description

Displays the description of the group.

Associated Rules

Displays the number of rules added to the group.

Associated Actions

Displays the number of group actions added to the group.

Applicable Clusters

Displays the number of clusters included in the group.

Applicable SVMs

Displays the number of SVMs included in the group.

Applicable Volumes

Displays the number of volumes included in the group.

Group Rules tab

The Group Rules tab displays the group rules you created for groups to contain storage objects. You can perform tasks such as adding, editing, or deleting a group rule. The tab also displays the group name for which the group rule is created and the storage object for which the rule is applied. You can also view the number of storage objects that satisfy the group rule.

Command buttons

Add

Displays the Add Group Rule dialog box, which enables you to create group rules for storage objects.

Edit

Displays the Edit Group Rule dialog box, which enables you to reconfigure previously configured group rules.

Delete

Deletes the selected group rule.

List view

The list view displays, in tabular format, the group rules you created for a specific storage object (clusters, volumes, or SVMs)and the count of storage objects that satisfy the defined group rule.

Name

Displays the name of the rule.

Associated Group

Displays the name of the group for which the group rule is defined.

Target Object Type

Displays the type of storage object to which the group rule is applied.

Applicable Objects

Displays the count of the storage objects included in the group based on the group rule.

Group Actions tab

The Group Actions tab displays the name and type of group actions you define for groups. You can perform tasks such as adding, editing, deleting, or reordering the group actions. The tab also displays the name of the group on which the group action is applied.

Command buttons

Add

Displays the Add Action dialog box, which enables you to create group actions for a group of storage objects. For example, you can set the threshold levels of storage objects in a group.

Edit

Displays the Edit Action dialog box, which enables you to reconfigure previously configured group actions.

Delete

Deletes the selected group action.

Reorder

Displays the Reorder Group Actions dialog box to rearrange the order of the group actions.

List view

The list view displays, in tabular format, the group actions you created for the groups in the Unified Manager server. You can use the column filters to customize the data that is displayed.

Rank

Displays the order of the group actions to be applied on the storage objects in a group.

Name

Displays the name of the group action.

Associated Group

Displays the name of the group for which the group action is defined.

Action Type

Displays the type of group action that you can perform on the storage objects in a group.

You cannot create multiple group actions of the same action type for a group. For example, you can create a group action of setting volume thresholds for a group. However, you cannot create another group action for the same group to change volume thresholds.

Description

Displays the description of the group action.

Related concepts

How group rules work for groups on page 260

Add Group dialog box

The Add Group dialog box enables you to create groups to include clusters, volumes, and Storage Virtual Machines (SVMs) based on the group rules.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Specifies the name of the group. You must enter a unique name for the group.

Description

Specifies a meaningful description of the group.

Command buttons

The command buttons enable you to add or cancel the creation of a new group.

Add

Creates the new group.

Cancel

Closes the Add Group dialog box without saving your changes.

Edit Group dialog box

The Edit Group dialog box enables you to change the name and description of a group.

You must have the OnCommand Administrator or Storage Administrator role.

Group Name

Displays the name of the group. When changing the group name, you must not use an existing group name.

Description

Provides a meaningful description of the group. You can edit this field when you want to change the current description of the group.

Command buttons

The command buttons enable you to save or cancel changes you make to the group.

Save

Saves the changes you made and closes the dialog box.

Cancel

Closes the Edit Group dialog box without saving your changes.

Groups details page

From the Groups details page, you can view the details of a selected group. You can also view additional information such as the group rules and group actions associated with the selected group.

- *Command buttons* on page 273
- *Summary area* on page 273

Command buttons

View Groups

Enables you to navigate to the Groups page.

Actions

Enables you to edit or delete the group, based on your role. You must have the OnCommand Administrator or Storage Administrator role.

Manage Group Rules

Enables you to navigate to the Group Rules page, which displays rules for this group.

Manage Group Actions

Enables you to navigate to the Group Actions page, which displays actions for this group.

Summary area

You can view the following group details:

Description

Displays the description provided for the group.

Created by

Displays the name of the user who created the group.

Creation Date

Displays the date when the group was created.

Associated Rules

Displays all the group rules created for a group, in tabular format. You can view the details of each group rule, such as the rule name, associated object type, and the count of storage objects of the associated object type.

Associated Actions

Displays all the group actions, configured for a group, in tabular format. You can view the details of each group action, such as the rank, name, action type, and description.

Add Group Rule dialog box

The Add Group Rule dialog box enables you to create group rules in Unified Manager to dynamically group storage objects. You can later configure and apply group actions for the group.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Specifies the name of the group rule.

Target Object Type

Specifies the type of storage objects to include in the group.

Group

Specifies the name of the group for which the group rule is created.

Conditions

Specifies conditions that determine which storage objects can be included in a group.

Condition group

Specifies condition groups which have one or more conditions defined for including storage objects in a group.

Command buttons

Save and Add

Adds the group rule and enables you to add another group rule without closing the dialog box.

Add

Adds the group rule and closes the Add Group Rule dialog box.

Cancel

Cancels the changes and closes the Add Group Rule dialog box.

Add Condition

Adds a condition to define the group rule.

Add Condition Group

Adds a condition group to define conditions for the group rule.

Related concepts

How group rules work for groups on page 260

Edit Group Rule dialog box

You can edit the group rules you created to include the maximum number of storage objects in a group.

You must have the OnCommand Administrator or Storage Administrator role.

Rule Name

Displays the name of the rule.

Target Object Type

Displays the storage object to be added to a selected group. You cannot change the object type.

Associated Group

Displays the associated group. You can select a different group for the group rule.

Condition

Displays the list of conditions for a selected group. You can edit the conditions. The storage objects are either removed or added to a selected group based on the changes.

Command buttons

Save

Saves the changes you made and closes the dialog box.

Cancel

Closes the Edit Group Rule dialog box without saving your changes.

Related concepts

How group rules work for groups on page 260

Add Group Action dialog box

The Add Group Action dialog box enables you to configure group actions that can be applied to storage objects of a selected group in Unified Manager.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Specifies the name of the action.

Description

Specifies the description of the action.

Group

Specifies the group for which the action is configured.

Action type

Specifies the type of action configured. Based on the selected action type, the Add Group Action dialog box expands, enabling you to configure a group action by providing the required values.

Unified Manager as of now supports only volume threshold action type.

Command buttons

Add

Adds the new action and closes the dialog box.

Cancel

Closes the Add Group Action dialog box dialog box without saving your changes.

Related tasks

Configuring volume thresholds for groups on page 268

Related references

Group action-volume thresholds section on page 275

Group action-volume thresholds section

The group action-volume thresholds section enables you to configure group-level thresholds for volumes. These thresholds are applied to all the volumes in a group. When the volume thresholds are configured at the group level, the global threshold values are not affected.

You can configure volume thresholds for the following to configure a group action:

- Capacity
- Growth
- Qtree quota
- Snapshot copies
- Inodes

Global default values are used if volume thresholds are not configured for any of these categories. You can set thresholds for the following:

- Capacity
- Growth
- Qtree quota
- Snapshot copies
- Inodes

Capacity section

You can set conditions for the following volume capacity thresholds:

Space Nearly Full

Specifies the percentage at which a volume is considered to be nearly full:

- Default value: 80 percent

  The value for this threshold must be lower than the value for the Volume Full threshold for the management server to generate an event.
- Event generated: Volume Nearly Full
- Event severity: Warning

Space Full

Specifies the percentage at which a volume is considered full:

- Default value: 90 percent
- Event generated: Volume Full
- Event severity: Error

Overcommitted

Specifies the percentage at which a volume is considered to be overcommitted:

- Default value: 100 percent
- Event generated: Volume Overcommitted
- Event severity: Error

Growth section

You can set the following threshold conditions for volume growth:

Growth Rate

Specifies the percentage at which a volume's growth rate is considered to be normal before the system generates a Volume Growth Rate Abnormal event:

- Default value: 1 percent
- Event generated: Volume Growth Rate Abnormal
- Event severity: Warning

Growth Rate Sensitivity

Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated.

A lower value for growth rate sensitivity indicates that the aggregate is highly sensitive to changes in the growth rate. The range for the growth rate sensitivity is 1 through 5.

- Default value: 2

Qtree Quota section

You can set the following volume quota threshold conditions:

Nearly Overcommitted

Specifies the percentage at which a volume is considered to be nearly overcommitted by qtree quotas:

- Default value: 95 percent
- Event generated: Volume Qtree Quota Nearly Overcommitted
- Event severity: Warning

Overcommitted

Specifies the percentage at which a volume is considered to be overcommitted by qtree quotas:

- Default value: 100 percent
- Event generated: Volume Qtree Quota Overcommitted
- Event severity: Error

Snapshot Copies section

You can set the following threshold conditions for the Snapshot copies in the volume:

Snapshot Reserve Full

Specifies the percentage at which the space reserved for Snapshot copies is considered full:

- Default value: 90 percent
- Event generated: Volume Snapshot Reserve Full
- Event severity: Error

Days Until Full

Specifies the number of days remaining before the space reserved for Snapshot copies reaches full capacity:

- Default value: 7
- Event generated: Volume Snapshot Reserve Days Until Full
- Event severity: Error

Count

Specifies the number of Snapshot copies that can be created on a volume before the system generates the Too Many Snapshot Copies event:

- Default value: 250
- Event generated: Too Many Snapshot Copies
- Event severity: Error

Inodes section

You can set the following threshold conditions for inodes:

Nearly Full

Specifies the percentage at which a volume is considered to have consumed most of its inodes:

- Default value: 80 percent
- Event generated: Inodes Nearly Full
- Event severity: Warning

Full

Specifies the percentage at which a volume is considered to have consumed all of its inodes:

- Default value: 90 percent
- Event generated: Inodes Full
- Event severity: Error

Edit Group Action dialog box

You can edit the group action that you created for groups by using the Edit Group Action dialog box.

You must have the OnCommand Administrator or Storage Administrator role.

Action Name

Displays the name of the group action.

Description

Displays the description of the group action.

Group

Displays the name of the group selected.

Action type

Displays the type of group action. You cannot change the action type. However, you can modify the parameters that you used to configure the group action.

Command buttons

Save

Saves the changes you made to the group action.

Cancel

Closes the Edit Group Action dialog box without saving your changes.

Reorder Group Actions dialog box

You can use the Reorder Group Actions dialog box to change the ranks of one or more group actions. The position of a group action in the grid determines the rank for the group action.

You must have the OnCommand Administrator or Storage Administrator role.

Rank

Specifies the order of the group action to be applied on storage objects in a group.

Name

Specifies the name of the group action.

Action Type

Specifies the type of action that you can perform on the storage objects in a group.

Associated Group

Specifies the name of the group for which the group actions are defined.

Related tasks

Reordering group actions on page 269

Managing and monitoring protection relationships

OnCommand Unified Manager enables you to create protection relationships, to monitor and troubleshoot SnapMirror and SnapVault relationships on managed clusters, and to restore data when it is overwritten or lost.

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

What resource pools are

Resource pools are groups of aggregates that are created by a storage administrator using Unified Manager to provide provisioning to partner applications for backup management.

You might pool your resources based on attributes such as performance, cost, physical location, or availability. By grouping related resources into a pool, you can treat the pool as a single unit for monitoring and provisioning. This simplifies the management of these resources and allows for a more flexible and efficient use of the storage.

Three variations of SnapMirror protection

Depending on the deployment of your data storage topology, Unified Manager enables you to configure three variations of SnapMirror protection relationships. All variations of SnapMirror protection offer failover disaster recovery protection, but offer differing capabilities in performance, version flexibility, and multiple backup copy protection.

Traditional SnapMirror protection relationships

Traditional SnapMirror protection provides block replication mirror protection between source and destination volumes.

In traditional SnapMirror relationships, mirror operations execute faster than they would in alternative SnapMirror relationships because the mirror operation is based on block replication; however, traditional SnapMirror protection requires that the destination volume run under the same or later version of Data ONTAP as the source volume.

SnapMirror protection with version-flexible replication

SnapMirror protection with version-flexible replication provides logical replication mirror protection between source and destination volumes, even if those volumes are running under different versions of Data ONTAP version 8.3 or higher.

In SnapMirror relationships with version-flexible replication, mirror operations do not execute as quickly as they would in traditional SnapMirror relationships.

Because of slower execution, SnapMirror with version-flexible replication protection is not suitable to implement in either of the following circumstances:

- The source object contains more than 10 million files to protect.
- The recovery point objective for the protected data is two hours or less. (That is, the destination must always contain mirrored, recoverable data that is no more than two hours older than data at the source.)

In either of the listed circumstances, the faster block-replication based execution of default SnapMirror protection is required.

SnapMirror protection with version-flexible replication and backup option

SnapMirror protection with version-flexible replication and version-flexible replication provides mirror protection between source and destination volumes and the capability to store multiple copies of the mirrored data at the destination.

The storage administrator can specify which Snapshot copies are mirrored from source to destination and can also specify how long to retain those copies at the destination, even if they are deleted at the source.

In SnapMirror relationships with version-flexible replication and version-flexible replication, mirror operations do not execute as quickly as they would in traditional SnapMirror relationships.

Viewing volume protection relationships

From the Volume Protection Relationships page, you can view existing volume SnapMirror and SnapVault relationships to monitor the status of and examine details about protection relationships, including transfer and lag status, source and destination details, schedule and policy information, and so on. You can also initiate relationship commands from the Volume Protection Relationships page.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Step

1. Click **Protection > Volume Relationships**.

Result

The Volume Protection Relationships page is displayed.

Related tasks

Adding a user on page 379

Creating a SnapVault protection relationship from the Volumes page

You can use the Volumes page to create SnapVault relationships for one or more volumes on the same Storage Virtual Machine (SVM) to enable data backups for protection purposes on volumes running Data ONTAP 8.2 or later.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Volumes** page, right-click a volume you want to protect and select **Protect**.

   Alternatively, to create multiple protection relationships on the same SVM, select one or more volumes in the Volumes page, and click **Protect** on the toolbar.

2. Select **SnapVault** from the menu.

   The Configure Protection dialog box is launched.

3. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary volume information.
4. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **SnapVault** tab.
6. Click **Apply**.

   You are returned to the Volumes page.

7. Click the protection configuration job link at the top of the **Volumes** page.

   If you are creating only one protection relationship, the Job details page is displayed; however, if you are creating more than one protection relationship, a filtered list of all the jobs associated with the protection operation is displayed.

8. Do one of the following:
   - If you have only one job, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
   - If you have more than one job:
     a. Click a job in the jobs list.
     b. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
     c. Use the **Back** button to return to the filtered list and view another job.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Related references

Configure Protection dialog box on page 337
Advanced Secondary Settings dialog box on page 326

Creating a SnapVault protection relationship from the Volume details page

You can create a SnapVault relationship using the Volume details page so that data backups are enabled for protection purposes on volumes running Data ONTAP 8.2 or later.

Before you begin

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

- You must have set up Workflow Automation to perform this task.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Protection** tab of the **Volume details** page, right-click a volume in the topology view that you want to protect.
2. Select **Protect > SnapVault** from the menu.

   The Configure Protection dialog box is launched.
3. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary resource information.
4. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.

   You are returned to the Volume details page.
7. Click the protection configuration job link at the top of the **Volume details** page.

   The Job details page is displayed.
8. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

   When the job tasks are complete, the new relationships are displayed in the Volume details page topology view.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Related references

Configure Protection dialog box on page 337
Advanced Secondary Settings dialog box on page 326

Creating a SnapMirror protection relationship from the Volume details page

You can use the Volume details page to create a SnapMirror relationship so that data replication is enabled for protection purposes. SnapMirror replication enables you to restore data from the destination volume in the event of data loss on the source for volumes running Data ONTAP 8.2 or later.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

Steps

1. In the **Protection** tab of the **Volume details** page, right-click in the topology view the name of a volume that you want to protect.
2. Select **Protect > SnapMirror** from the menu.

   The Configure Protection dialog box is displayed.
3. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
4. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.

   You are returned to the Volume details page.
7. Click the protection configuration job link at the top of the **Volume details** page.

   The job's tasks and details are displayed in the Job details page.
8. In the **Job details** page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
9. When the job tasks are complete, click **Back** on your browser to return to the **Volume details** page.

   The new relationship is displayed in the Volume details page topology view.

Result

Depending on the destination SVM you specified during configuration or on the options you enabled in your Advanced settings, the resulting SnapMirror relationship might be one of several possible variations:

- If you specified a destination SVM that runs under the same or a newer version of Data ONTAP compared to that of the source volume, a block-replication-based SnapMirror relationship is the default result.

- If you specified a destination SVM that runs under the same or a newer version of Data ONTAP (8.3 or higher) compared to that of the source volume, but you enabled version-flexible replication in the Advanced settings, a SnapMirror relationship with version-flexible replication is the result.
- If you specified a destination SVM that runs under an earlier version of Data ONTAP 8.3, or a version that is higher than that of the source volume and the earlier version supports version-flexible replication, a SnapMirror relationship with version-flexible replication is the automatic result.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Related references

Configure Protection dialog box on page 337
Advanced Destination Settings dialog box on page 327

Creating a SnapMirror protection relationship from the Volumes page

Using the Volumes page enables you to create several SnapMirror protection relationships at one time by selecting more than one volume on the same Storage Virtual Machine (SVM). The volumes must be running Data ONTAP 8.2 or later.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Volumes** page, locate a volume that you want to protect and right-click it.

   Alternatively, to create multiple protection relationships on the same SVM, select one or more volumes in the Volumes page, and click **Protect > SnapMirror** on the toolbar.

   The Configure Protection dialog box is displayed.
2. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
3. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
4. Complete the **Destination Information** area and the **Relationship Settings** area in the **SnapMirror** tab.
5. Click **Apply**.

   You are returned to the Volumes page.

6. Click the protection configuration job link at the top of the **Volumes** page.

   If you are creating only one protection relationship, the Job details page is displayed; however, if you are creating more than one protection relationship, a list of all the jobs associated with the protection operation is displayed.

7. Do one of the following:

   - If you have only one job, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
   - If you have more than one job:
     a. Click a job in the jobs list.
     b. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
     c. Use the **Back** button to return to the filtered list and view another job.

Result

Depending on the destination SVM you specified during configuration or on the options you enabled in your Advanced settings, the resulting SnapMirror relationship might be one of several possible variations:

- If you specified a destination SVM that runs under the same or a newer version of Data ONTAP compared to that of the source volume, a block-replication-based SnapMirror relationship is the default result.
- If you specified a destination SVM that runs under the same or a newer version of Data ONTAP (8.3 or higher) compared to that of the source volume, but you enabled version-flexible replication in the Advanced settings, a SnapMirror relationship with version-flexible replication is the result.
- If you specified a destination SVM that runs under an earlier version of Data ONTAP 8.3 or a later version than that of the source volume, and the earlier version supports version-flexible replication, a SnapMirror relationship with version-flexible replication is the automatic result.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Related references

Configure Protection dialog box on page 337
Advanced Destination Settings dialog box on page 327

Creating a SnapMirror relationship with version-flexible replication

You can create a SnapMirror relationship with version-flexible replication. Version-flexible replication enables you to implement SnapMirror protection even if source and destination volumes run under different versions of Data ONTAP 8.3 or higher.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

- The source and destination SVMs must each have a SnapMirror license enabled.
- The source and destination SVMs must each run under a version of Data ONTAP (8.3 or higher) that supports version-flexible replication.

About this task

SnapMirror with version-flexible replication enables you to implement SnapMirror protection even in heterogeneous storage environments in which not all storage is running under one version of Data ONTAP; however, mirror operations performed under SnapMirror with version-flexible replication do not execute as quickly as they would under traditional block replication SnapMirror.

Steps

1. Display the **Configure Protection** dialog box for the volume that you want to protect.
    - If you are viewing the Protection tab of the Volume details page, right-click in the topology view that has the name of a volume that you want to protect and select **Protect > SnapMirror** from the menu.
    - If you are viewing the Volumes page, locate a volume that you want to protect and right-click it; then select **Protect > SnapMirror** from the menu.

    The Configure Protection dialog box is displayed.
2. Click **SnapMirror** to view the **SnapMirror** tab.
3. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.

    If you specify a destination SVM that runs under an earlier version of Data ONTAP than the source volume you are protecting, and if that earlier version supports version-flexible replication, this task automatically configures SnapMirror with version-flexible replication.
4. If you specify a destination SVM that runs under the same version of Data ONTAP as that of the source volume, but you still want to configure SnapMirror with version-flexible replication, click **Advanced** to enable version-flexible replication and then click **Apply**.
5. Click **Apply**.

    You are returned to the Volume details page.
6. Click the protection configuration job link at the top of the **Volume details** page.

    The jobs tasks and details are displayed in the Job details page.
7. In the **Job details** page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
8. When the job tasks are complete, click **Back** on your browser to return to the **Volume details** page.

    The new relationship is displayed in the Volume details page topology view.

Creating SnapMirror relationships with version-flexible replication with backup option

You can create a SnapMirror relationship with version-flexible replication and backup option capability. Backup option capability enables you to implement SnapMirror protection and also retain multiple versions of backup copies at the destination location.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- The source and destination SVMs must each have a SnapMirror license enabled.
- The source and destination SVMs must each have a SnapVault license enabled.
- The source and destination SVMs must each run under a version of Data ONTAP (8.3 or higher) that supports version-flexible replication.

About this task

Configuring SnapMirror with backup option capability enables you to protect your data with SnapMirror disaster recovery capabilities, such as volume failover ability, and at the same time provide SnapVault capabilities, such as multiple backup copy protection.

Steps

1. Display the **Configure Protection** dialog box for the volume that you want to protect.
    - If you are viewing the Protection tab of the Volume details page, right-click in the topology view the name of a volume that you want to protect and select **Protect > SnapMirror** from the menu.
    - If you are viewing the Volumes page, locate a volume you want to protect and right-click it; then select **Protect > SnapMirror** from the menu.

    The Configure Protection dialog box is displayed.
2. Click **SnapMirror** to view the **SnapMirror** tab.
3. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
4. Click **Advanced** to display the **Advanced Destination Settings** dialog box.
5. If the **Version-Flexible Replication** check box is not already selected, select it now.
6. Select the **With backup option** check box to enable backup option capability; then click **Apply**.
7. Click **Apply**.

    You are returned to the Volume details page.
8. Click the protection configuration job link at the top of the **Volume details** page.

    The jobs tasks and details are displayed in the Job details page.
9. In the **Job details** page, click **Refresh** to update the task list and task details associated with the protection configuration job and to determine when the job is complete.

10. When the job tasks are complete, click **Back** on your browser to return to the **Volume details** page.

    The new relationship is displayed in the Volume details page topology view.

Configuring destination efficiency settings

You can configure destination efficiency settings such as deduplication, compression, autogrow, and space guarantee on a protection destination using the Advanced Destination Settings dialog box. You use these settings when you want to maximize space utilization on a destination or secondary volume.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

By default, efficiency settings match those of the source volume, except for compression settings in a SnapVault relationship, which are disabled by default.

Steps

1. Click either the **SnapMirror** tab or the **SnapVault** tab in the **Configure Protection** dialog box, depending on the type of relationship you are configuring.
2. Click **Advanced** in the **Destination Information** area.

   The Advanced Destination Settings dialog box is opened.
3. Enable or disable the efficiency settings for deduplication, compression, autogrow, and space guarantee, as required.
4. Click **Apply** to save your selections and return to the **Configure Protection** dialog box.

Creating SnapMirror and SnapVault schedules

You can create basic or advanced SnapMirror and SnapVault schedules to enable automatic data protection transfers on a source or primary volume so that transfers take place more frequently or less frequently, depending on how often the data changes on your volumes.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role..
- You must have already completed the Destination Information area in the Configure Protection dialog box.
- You must have set up Workflow Automation to perform this task.

Steps

1. From the **SnapMirror** tab or **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Schedule** link in the **Relationship Settings** area.

   The Create Schedule dialog box is displayed.
2. In the **Schedule Name** field, type the name you want to give to the schedule.

3. Select one of the following:
    - **Basic**
      Select if you want to create a basic interval-style schedule.
    - **Advanced**
      Select if you want to create a cron-style schedule.
4. Click **Create**.

    The new schedule is displayed in the SnapMirror Schedule or SnapVault Schedule drop-down list.

Related references

Create Schedule dialog box on page 341

Creating cascade or fanout relationships to extend protection from an existing protection relationship

You can extend protection from an existing relationship by creating either a fanout from the source volume or a cascade from the destination volume of an existing relationship. You might do this when you need to copy data from one site to many sites or to provide additional protection by creating more backups.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Protection > Volume Relationships**.
2. From the **Volume Protection Relationships** page, select the SnapMirror relationship from which you want to extend protection.
3. On the action bar, click **Extend Protection**.
4. In the menu, select either **From Source** or **From Destination**, depending on whether you are creating a fanout relationship from the source or a cascade relationship from the destination.
5. Select either **With SnapMirror** or **With SnapVault**, depending on the type of protection relationship you are creating.

    The Configure Protection dialog box is displayed.
6. Complete the information as indicated in the **Configure Protection** dialog box.

Related references

Configure Protection dialog box on page 337

Editing protection relationships

You can edit existing protection relationships to change the maximum transfer rate, the protection policy, or the protection schedule. You might edit a relationship to decrease the bandwidth used for transfers, or to increase the frequency of scheduled transfers because data is changing often.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must be running Data ONTAP 8.2 or later.

About this task

The selected volumes must be protection relationship destinations. You cannot edit relationships when source volumes, load-sharing volumes, or volumes that are not the destination of a SnapMirror or SnapVault relationship are selected.

Steps

1. From the **Volume Protection Relationships** page, select in the volumes list one or more volumes in the same Storage Virtual Machine (SVM) for which you want to edit relationship settings, and then select **Edit** from the toolbar.

   The Edit Relationship dialog box is displayed.

2. In the **Edit Relationship** dialog box, edit the maximum transfer rate, protection policy, or protection schedule, as needed.
3. Click **Apply**.

   The changes are applied to the selected relationships.

Editing protection relationships from the Volume details page

You can edit existing protection relationships to change the current maximum transfer rate, protection policy, or protection schedule. You might edit a relationship to decrease the bandwidth used for transfers, or to increase the frequency of scheduled transfers because data is changing often.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must be running Data ONTAP 8.2 or later.
- You must have installed and configured Workflow Automation.

About this task

The selected volumes must be protection relationship destinations. You cannot edit relationships when source volumes, load-sharing volumes, or volumes that are not the destination of a SnapMirror or SnapVault relationship are selected.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the protection relationship you want to edit and right-click it.
2. Select **Edit** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship > Edit** to edit the relationship for which you are currently viewing the details.

   The Edit Relationship dialog box is displayed.
3. In the **Edit Relationship** dialog box, edit the maximum transfer rate, protection policy, or protection schedule, as needed.
4. Click **Apply**.

   The changes are applied to the selected relationships.

Creating SnapMirror policies

You can create a SnapMirror policy to specify the SnapMirror transfer priority for protection relationships. SnapMirror policies enable you to maximize transfer efficiency from the source to the destination by assigning priorities so that lower-priority transfers are scheduled to run after normal-priority transfers.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation to enable this operation.
- This task assumes that you have already completed the Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapMirror** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

   The Create SnapMirror Policy dialog box is displayed.
2. In the **Policy Name** field, type a name you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority you want to assign to the policy.
4. In the **Comment** field, enter an optional comment for the policy.
5. Click **Create**.

   The new policy is displayed in the SnapMirror Policy drop-down list.

Related references

Create SnapMirror Policy dialog box on page 342

Creating SnapVault policies

You can create a new SnapVault policy to set the priority for a SnapVault transfer. You use policies to maximize the efficiency of transfers from the primary to the secondary in a protection relationship.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation to enable this operation.
- You must have already completed Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

   The SnapVault tab is displayed.
2. In the **Policy Name** field, type the name that you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority that you want to assign to the policy.
4. Optional: In the **Comment** field, enter a comment for the policy.
5. In the **Replication Label** area, add or edit a replication label, as necessary.
6. Click **Create**.

   The new policy is displayed in the Create Policy drop-down list.

Related references

Create SnapVault Policy dialog box on page 343

Aborting an active data protection transfer from the Volume details page

You can abort an active data protection transfer when you want to stop a SnapMirror replication that is in progress. You can also clear the restart checkpoint for a transfer if it is not a baseline transfer. You might abort a transfer when it conflicts with another operation, such as a volume move.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role..
- You must have set up Workflow Automation to perform the operation.

About this task

The abort action does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

- When the volume is a Data ONTAP 8.1 or earlier cluster volume

You cannot clear the restart checkpoint for a baseline transfer.

Steps

1. In the **Protection** tab of the **Volume details** page, right-click the relationship in the topology view for the data transfer you want to abort and select **Abort**.

   The Abort Transfer dialog box is displayed.

2. If you want to clear the restart checkpoint for a transfer that is not a baseline transfer, select **Clear Checkpoints**.
3. Click **Continue**.

   The Abort Transfer dialog box is closed, and the status of the abort operation displays at the top of the Volume details page along with a link to the job details.

4. Optional: Click the **View details** link to go to the **Job details** page for additional details and to view job progress.
5. Click each job task to view its details.
6. Click the Back arrow on your browser to return to the **Volume details** page.

   The abort operation is finished when all job tasks successfully complete.

Aborting an active data protection transfer

You can abort an active data protection transfer when you want to stop a SnapMirror replication that is in progress. You can also clear the restart checkpoint for transfers subsequent to the baseline transfer. You might abort a transfer when it conflicts with another operation, such as a volume move.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role..
- You must have set up Workflow Automation to perform the operation.

About this task

The abort action does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 or earlier cluster volume

You cannot clear the restart checkpoint for a baseline transfer.

Steps

1. To abort transfers for one or more protection relationships, from the **Volume Protection Relationships** page, select one or more volumes and, on the toolbar, click **Abort**.

   The Abort Transfer dialog box is displayed.

2. If you want to clear the restart checkpoint for a transfer that is not a baseline transfer, select **Clear Checkpoints**.

3. Click **Continue**.

   The Abort Transfer dialog box is closed, and the status of the abort job displays at the top of the Volume Protection Relationships page, along with a link to the job details.

4. Optional: Click the **View details** link to go to the **Job details** page for additional details and to view job progress.

Quiescing protection relationships from the Volume details page

You can quiesce a protection relationship to temporarily prevent data transfers from occurring. You might quiesce a relationship when you want to create a Snapshot copy of a SnapMirror destination volume that contains a database, and you want to ensure that its contents are stable during the Snapshot copy.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation to perform this task.

About this task

The quiesce action does not display in the following instances:

- If RBAC settings do not allow this action, for example, if you have only operator privileges
- When the volume ID is unknown, for example, when you have a intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 or earlier cluster volume
- When you have not paired Workflow Automation and Unified Manager

Steps

1. In the **Protection** tab of the **Volume details** page, right-click the relationship in the topology view for the protection relationship that you want to quiesce.
2. Select **Quiesce** from the menu.
3. Click **Yes** to continue.

   The status of the quiesce job is displayed at the top of the Volume details page, along with a link to the job details.

4. Click the **View details** link to go to the **Job details** page for additional details and job progress.
5. Optional: Click the Back arrow on your browser to return to the **Volume details** page.

   The quiesce job is finished when all job tasks are successfully completed.

Quiescing a protection relationship

From the Volume Protection Relationships page, you can quiesce a protection relationship to temporarily prevent data transfers from occurring. You might quiesce a relationship when you want to create a Snapshot copy of a SnapMirror destination volume that contains a database, and you want to ensure that its contents are stable during the Snapshot copy operation.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

About this task

The quiesce action does not display in the following instances:

- If RBAC settings do not allow this action; for example, if you have only operator privileges
- When the volume ID is unknown; for example, when you have a intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 or earlier cluster volume
- When you have not paired Workflow Automation and Unified Manager

Steps

1. To quiesce transfers for one or more protection relationships, from the **Volume Protection Relationships** page, select one or more volumes and, on the toolbar, click **Quiesce**.

   The Quiesce dialog box is displayed.
2. Click **Continue**.

   The status of the quiesce job is displayed at the top of the Volume details page, along with a link to the job details.
3. Click the **View details** link to go to the **Job details** page for additional details and job progress.
4. Optional: Click the **Back** arrow on your browser to return to the **Volume Protection Relationships** page.

   The quiesce job is finished when all job tasks are successfully completed.

Breaking a SnapMirror relationship from the Volume details page

You can break a protection relationship from the Volume details page and stop data transfers between a source and destination volume in a SnapMirror relationship. You might break a relationship when you want to migrate data, for disaster recovery, or for application testing. The destination volume is changed to a read-write volume. You cannot break a SnapVault relationship.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation to perform this task.

Steps

1. In the **Protection** tab of the **Volume details** page, select from the topology the SnapMirror relationship you want to break.
2. Right-click the destination and select **Break** from the menu.

The Break Relationship dialog box is displayed.

3. Click **Continue** to break the relationship.
4. In the topology, verify that the relationship is broken.

Breaking a SnapMirror relationship

You can break a protection relationship to stop data transfers between a source volume and a destination volume in a SnapMirror relationship. You might break a relationship when you want to migrate data, for disaster recovery, or for application testing. The destination volume is changed to a read/write volume. You cannot break a SnapVault relationship.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation to perform this task.

Steps

1. From the **Volume Protection Relationships** page, select one or more volumes with protection relationships for which you want to stop data transfers and, on the toolbar, click **Break**.

   The Break Relationship dialog box is displayed.

2. Click **Continue** to break the relationship.
3. In the **Volume Protection Relationships** page, verify in the **Relationship State** column that the relationship is broken.

   The Relationship State column is hidden by default, so you might need to select it in the show/hide column list ▥.

Removing a protection relationship from the Volume details page

You can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation to perform this task.

Steps

1. In the **Protection** tab of the **Volume details** page, select from the topology the SnapMirror relationship you want to remove.
2. Right-click the name of the destination and select **Remove** from the menu.

   The Remove Relationship dialog box is displayed.

3. Click **Continue** to remove the relationship.

   The relationship is removed from the Volume details page.

Removing a protection relationship

From the Volume Protection Relationships page, you can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. From the **Volume Protection Relationships** page, select one or more volumes with protection relationships you want to remove and, on the toolbar, click **Remove**.

   The Remove Relationship dialog box is displayed.

2. Click **Continue** to remove the relationship.

   The relationship is removed from the Volume Protection Relationships page.

Resuming scheduled transfers on a quiesced relationship

After you have quiesced a relationship to stop scheduled transfers from occurring, you can use **Resume** to reenable scheduled transfers so that data on the source or primary volume is protected. Transfers resume from a checkpoint, if one exists, at the next scheduled transfer interval.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

About this task

You can select no more than 10 quiesced relationships on which to resume transfers.

Steps

1. From the **Volume Protection Relationships** page, select one or more volumes with quiesced relationships, and, on the toolbar, click **Resume**.
2. In the **Resume** dialog box, click **Continue**.

   You are returned to the Volume Protection Relationships page.

3. To view the related job tasks and to track their progress, click the job link that is displayed at the top of the **Volume Protection Relationships** page.
4. Do one of the following:
   - If only one job is displayed, in the Job details page click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
   - If more than one job is displayed,

   a. In the Jobs page, click the job for which you want to view the details.
   b. In the Job details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

   After the jobs finish, data transfers are resumed at the next scheduled transfer interval.

Resuming scheduled transfers on a quiesced relationship from the Volume details page

After you have quiesced a relationship to stop scheduled transfers from occurring, you can use **Resume** on the Volume details page to reenable scheduled transfers so that data on the source or primary volume is protected. Transfers resume from a checkpoint, if one exists, at the next scheduled transfer interval.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation to perform this task.

Steps

1. In the **Protection** tab of the **Volume details** page, right-click in the topology view a quiesced relationship that you want to resume.

   Alternatively, select **Resume** from the **Actions > Relationship** menu.
2. In the **Resume** dialog box, click **Continue**.

   You are returned to the Volume details page.
3. To view the related job tasks and to track their progress, click the job link that is displayed at the top of the **Volume details** page.
4. In the **Job details** page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

   After the jobs are complete, data transfers are resumed at the next scheduled transfer interval.

Initializing or updating protection relationships from the Volume details page

You can perform a first-time baseline transfer on a new protection relationship, or update a relationship if it is already initialized and you want to perform a manual, unscheduled incremental update to transfer data immediately.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must be running Data ONTAP 8.2 or later.
- You must have set up OnCommand Workflow Automation to enable this operation.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the protection relationship that you want to initialize or update, and then right-click it.
2. Select **Initialize/Update** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship > Initialize/Update** to initialize or update the relationship for which you are currently viewing the details.

   The Initialize/Update dialog box is displayed.
3. In the **Transfer Options** tab, select a transfer priority and the maximum transfer rate.
4. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.
5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
6. Click **Submit**.

   You are returned to the Initialize/Update dialog box.
7. If you selected more than one source to initialize or update, click **Default** for the next read/write source for which you want to specify an existing Snapshot copy.

   You cannot select a different Snapshot copy for data protection volumes.
8. Click **Submit** to begin the initialization or update job.

   The initialization or update job is started, you are returned to the Volume details page, and a jobs link is displayed at the top of the page.
9. Optional: Click **View Jobs** on the **Volume details** page to track the status of each initialization or update job.

   A filtered list of jobs is displayed.
10. Optional: Click each job to see its details.
11. Optional: Click the Back arrow on your browser to return to the **Volume details** page.

    The initialization or update operation is finished when all job tasks successfully complete.

Initializing or updating protection relationships

From the Volume Protection Relationships page, you can perform a first-time baseline transfer on a new protection relationship, or update a relationship if it is already initialized and you want to perform a manual, unscheduled incremental update to transfer immediately.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must be running Data ONTAP 8.2 or later.
- You must have set up OnCommand Workflow Automation.

Steps

1. From the **Volume Protection Relationships** page, right-click a volume and select one or more volumes with relationships that you want to update or initialize, and then, on the toolbar, click **Initialize/Update**.

The Initialize/Update dialog box is displayed.

2. In the **Transfer Options** tab, select a transfer priority and the maximum transfer rate.
3. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.
4. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
5. Click **Submit**.

   You are returned to the Initialize/Update dialog box.
6. If you selected more than one source to initialize or update, click **Default** for the next source for which you want to specify an existing Snapshot copy.
7. Click **Submit** to begin the initialization or update job.

   The initialization or update job is started, you are returned to the Volume Protection Relationships page, and a jobs link is displayed at the top of the page.
8. Optional: Click **View Jobs** on the **Volumes** page to track the status of each initialization or update job.

   A filtered list of jobs is displayed.
9. Optional: Click each job to see its details.
10. Optional: Click the **Back** arrow on your browser to return to the **Volume Protection Relationships** page.

    The initialization or update operation is finished when all tasks successfully finish.

Resynchronizing protection relationships from the Volume details page

You can resynchronize data on a SnapMirror or SnapVault relationship that was broken and then the destination was made read/write so that data on the source matches the data on the destination. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must be running Data ONTAP 8.2 or later.
- You must have set up OnCommand Workflow Automation to perform this operation.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the protection relationship that you want to resynchronize and right-click it.
2. Select **Resynchronize** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship > Resynchronize** to resynchronize the relationship for which you are currently viewing the details.

   The Resynchronize dialog box is displayed.

3. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
4. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.
5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
6. Click **Submit**.

   You are returned to the Resynchronize dialog box.
7. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
8. Click **Submit** to begin the resynchronization job.

   The resynchronization job is started, you are returned to the Volume details page and a jobs link is displayed at the top of the page.
9. Optional: Click **View Jobs** on the **Volume details** page to track the status of each resynchronization job.

   A filtered list of jobs is displayed.
10. Optional: Click the Back arrow on your browser to return to the **Volume details** page.

    The resynchronization job is finished when all job tasks successfully complete.

Resynchronizing protection relationships

From the Volume Protection Relationships page, you can resynchronize a relationship either to recover from an event that disabled your source volume or when you want to change the current source to a different volume.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must be running Data ONTAP 8.2 or later.
- You must have set up Workflow Automation.

Steps

1. From the **Volume Protection Relationships** page, select one or more volumes with quiesced relationships and, from the toolbar, click **Resynchronize**.

   The Resynchronize dialog box is displayed.
2. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
3. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.
4. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
5. Click **Submit**.

   You are returned to the Resynchronize dialog box.

6. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
7. Click **Submit** to begin the resynchronization job.

   The resynchronization job is started, you are returned to the Volume Protection Relationships page, and a jobs link is displayed at the top of the page.
8. Optional: Click **View Jobs** on the **Volume Protection Relationships** page to track the status of each resynchronization job.

   A filtered list of jobs is displayed.
9. Optional: Click the **Back** arrow on your browser to return to the **Volume Protection Relationships** page.

   The resynchronization operation is finished when all job tasks successfully finish.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Reversing protection relationships from the Volume details page

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- Your system must be running Data ONTAP 8.2 or later.
- Workflow Automation must be set up to perform this operation.
- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.
- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on the reverse resynchronization relationship are the same as those on the original protection relationship.
  If policies and schedules do not exist, they are created.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the SnapMirror relationship on which you want to reverse the source and destination, and right-click it.
2. Select **Reverse Resync** from the menu.

   The Reverse Resync dialog box is displayed.
3. Verify that the relationship displayed in the **Reverse Resync** dialog box is the one for which you want to perform the reverse resynchronization operation, and then click **Submit**.

   The Reverse Resync dialog box is closed and a job link is displayed at the top of the Volume details page.
4. Optional: Click **View Jobs** on the **Volume details** page to track the status of each reverse resynchronization job.

   A filtered list of jobs is displayed.
5. Optional: Click the Back arrow on your browser to return to the **Volume details** page.

   The reverse resynchronization operation is finished when all job tasks are completed successfully.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Related references

Reverse Resync dialog box on page 350

Reversing protection relationships

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to a read/write volume while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must be running Data ONTAP 8.2 or later.
- You must have set up Workflow Automation.
- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.
- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.

- Policies and schedules created on reverse resynchronization relationships are the same as those on the original protection relationship.
  If policies and schedules do not exist, they are created.

Steps

1. From the **Volume Protection Relationships** page, select one or more volumes with relationships that you want to reverse, and, on the toolbar, click **Reverse Resync**.

   The Reverse Resync dialog box is displayed.
2. Verify that the relationships displayed in the **Reverse Resync** dialog box are the ones for which you want to perform the reverse resynchronization operation, and then click **Submit**.

   The reverse resynchronization operation is started, you are returned to the Volume Protection Relationships page, and a jobs link is displayed at the top of the page.
3. Optional: Click **View Jobs** on the **Volume Protection Relationships** page to track the status of each reverse resynchronization job.

   A filtered list of jobs related to this operation is displayed.
4. Optional: Click the **Back** arrow on your browser to return to the **Volume Protection Relationships** page.

   The reverse resynchronization operation is finished when all job tasks successfully complete.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Related references

Reverse Resync dialog box on page 350

Restoring data using the Volumes page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Volumes page. Restoring data from volumes that have a version of Data ONTAP that is earlier than 8.2 is not supported.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot restore NTFS file streams.

The restore option is not available in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Volumes** page, select a volume from which you want to restore data.
2. From the toolbar, click **Restore**.

The Restore dialog box is displayed.

3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
4. Select the items you want to restore.

   You can restore the entire volume, or you can specify folders and files you want to restore.
5. Select the location to which you want the selected items restored; either **Original Location** or **Alternate Location**.
6. Click **Restore**.

   The restore process begins.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Related references

Restore dialog box on page 328

Restoring data using the Volume details page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Volume details page. Restoring data from volumes that have a version of Data ONTAP that is earlier than 8.2 is not supported.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot restore NTFS file streams.

The restore option is not available in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Protection** tab of the **Volume details** page, right-click in the topology view the name of the volume that you want to restore.
2. Select **Restore** from the menu.

   Alternatively, select **Restore** from the **Actions** menu to protect the current volume for which you are viewing the details.

   The Restore dialog box is displayed.
3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
4. Select the items you want to restore.

   You can restore the entire volume, or you can specify folders and files you want to restore.

5. Select the location to which you want the selected items restored: either **Original Location** or **Alternate Existing Location**.
6. If you select an alternate existing location, do one of the following:
    - In the Restore Path text field, type the path of the location to which you want to restore the data and then click **Select Directory**.
    - Click **Browse** to launch the Browse Directories dialog box and complete the following steps:
        a. Select the cluster, SVM, and volume to which you want to restore.
        b. In the Name table, select a directory name.
        c. Click **Select Directory**.
7. Click **Restore**.

    The restore process begins.

Related concepts

Executing protection workflows using OnCommand Workflow Automation on page 357

Related references

Restore dialog box on page 328

Creating resource pools

You can use the Create Resource Pool dialog box to group aggregates for provisioning purposes.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Resource pools can contain aggregates from different clusters, but the same aggregate cannot belong to different resource pools.

Steps

1. Click **Storage > Resource Pools**.
2. On the **Resource Pools** page, click **Create**.

    The Create Resource Pool dialog box is displayed.
3. Follow the instructions in the dialog box to provide a name and description and to add aggregates as members to the resource pool you want to create.

Related tasks

Adding a user on page 379

Related references

Create Resource Pool dialog box on page 319

Editing resource pools

You can edit an existing resource pool when you want to change the resource pool name and the description.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The **Edit** button is enabled only when one resource pool is selected. If more than one resource pool is selected, the **Edit** button is disabled.

Steps

1. Click **Storage > Resource Pools**.
2. Select one resource pool from the list.
3. Click **Edit**.

   The Edit Resource Pool window is displayed.
4. Edit the resource pool name and description as needed.
5. Click **Save**.

   The new name and description are displayed in the resource pool list.

Related tasks

Adding a user on page 379

Viewing resource pools inventory

You can use the Resource Pools page to view the resource pool inventory and to monitor the remaining capacity for each resource pool.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Step

1. Click **Storage > Resource Pools**.

   The resource pool inventory is displayed.

Related tasks

Adding a user on page 379

Related references

Resource Pools page on page 316

Adding resource pool members

A resource pool consists of a number of member aggregates. You can add aggregates to existing resource pools to increase the amount of space available for secondary volume provisioning.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can add no more than 200 aggregates to a resource pool at one time. Aggregates shown in the Aggregates dialog box do not belong to any other resource pool.

Steps

1. Click **Storage > Resource Pools**.
2. Select a resource pool from the **Resource Pools** list.

   The resource pool members are displayed in the area below the resource pool list.
3. In the resource pool member area, click **Add**.

   The Aggregates dialog box is displayed.
4. Select one or more aggregates.
5. Click **Add**.

   The dialog box is closed and the aggregates are displayed in the member list for the selected resource pool.

Related tasks

Adding a user on page 379

Removing aggregates from resource pools

You can remove aggregates from an existing resource pool: for example, when you want to use an aggregate for some other purpose.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Resource pool members are displayed only when a resource pool is selected.

Steps

1. Click **Storage > Resource Pools**.
2. Select the resource pool from which you want to remove member aggregates.

   The list of member aggregates is displayed in the Members pane.
3. Select one or more aggregates.

The **Remove** button is enabled.

4. Click **Remove.**

   A warning dialog box is displayed.

5. Click **Yes** to continue.

   The selected aggregates are removed from the Members pane.

Related tasks

Adding a user on page 379

Deleting resource pools

You can delete resource pools when they are no longer needed. For example, you might want to redistribute the member aggregates from one resource pool to several other resource pools, making the original resource pool obsolete.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The **Delete** button is enabled only when at least one resource pool is selected.

Steps

1. Click **Storage** > **Resource Pools**.
2. Select the resource pool you want to delete.
3. Click **Delete**.

   The resource pool is removed from the resource pool list and its aggregates are removed from the members list.

Related tasks

Adding a user on page 379

Related references

Resource Pools page on page 316

Understanding SVM associations

Storage Virtual Machine (SVM) associations are mappings from a source SVM to a destination SVM that are used by partner applications for resource selection and secondary volume provisioning.

Associations are always created between a source SVM and a destination SVM, regardless of whether the destination SVM is a secondary destination or a tertiary destination. You cannot use a secondary destination SVM as a source to create an association with a tertiary destination SVM.

You can associate SVMs in three ways:

- Associate any SVM

You can create an association between any primary source SVM and one or more destination SVMs. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVMs. For example, you might want applications from several different sources at different locations to be backed up to one or more destination SVMs in one location.

- Associate a particular SVM
  You can create an association between a specific source SVM and one or more specific destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, you can choose this option to associate a specific source SVM to a specific destination SVM that is assigned to only that client.
- Associate with an external SVM
  You can create an association between a source SVM and an external flexible volume of a destination SVM.

Storage Virtual Machine (SVM) and resource pool requirements to support storage services

You can better ensure conformance in partner applications if you observe some SVM association and resource pool requirements that are specific to storage services: for example, when you associate Storage Virtual Machines (SVMs) and create resource pools in Unified Manager to support a protection topology in a storage service provided by a partner application.

Some applications partner with the Unified Manager server to provide services that automatically configure and execute SnapMirror or SnapVault backup protection between source volumes and protection volumes in secondary or tertiary locations. To support these protection storage services, you must use Unified Manager to configure the necessary SVM associations and resource pools.

To support storage service single-hop or cascaded protection, including replication from a SnapMirror source or SnapVault primary volume to either destination SnapMirror or to SnapVault backup volumes that reside in secondary or tertiary locations, observe the following requirements:

- SVM associations must be configured between the SVM containing the SnapMirror source or SnapVault primary volume and any SVM on which either a secondary volume or a tertiary volume resides.
    - For example, to support a protection topology in which source volume Vol_A resides on SVM_1, and SnapMirror secondary destination volume Vol_B resides on SVM_2, and tertiary SnapVault backup volume Vol_ C resides on SVM_3, you must use the Unified Manager web UI to configure a SnapMirror association between SVM_1 and SVM_2 and a SnapVault backup association between SVM_1 and SVM_3.
      In this example, any SnapMirror association or SnapVault backup association between SVM_2 and SVM_3 is not necessary and is not used.
    - To support a protection topology in which both source volume Vol_A and SnapMirror destination volume Vol_B reside on SVM_1, you must configure a SnapMirror association between SVM_1 and Vserver_1.
- The resource pools must include cluster aggregate resources available to the associated SVMs.
  You configure resource pools through the Unified Manager web UI and then assign through the partner application the storage service secondary target and tertiary target nodes.

Related tasks

Creating resource pools on page 307
Creating Storage Virtual Machine (SVM) associations on page 312

Creating Storage Virtual Machine (SVM) associations

The Create Storage Virtual Machine Associations wizard enables partner protection applications to associate a source Storage Virtual Machine (SVM) with a destination SVM for use with SnapMirror and SnapVault relationships. Partner applications use these associations at the time of initial provisioning of destination volumes to determine which resources to select.

Before you begin

- The SVM you are associating must already exist.
- You must have the OnCommand Administrator or Storage Administrator role.

About this task

For any source SVM and relationship type, you can choose only one destination SVM on each destination cluster.

Changing associations using the delete and create functions affects only future provisioning operations. It does not move existing destination volumes.

Steps

1. Click **Storage > Storage Virtual Machine Associations**.
2. Click **Create**.

   The Create Storage Virtual Machine Associations wizard is launched.
3. Select one of the following sources:
   - **Any**
     Choose this option when you want to create an association between any primary SVM source to one or more destination SVM. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVM. For example, you might want applications from several different sources at different locations backed up to one or more destination SVM in one location.
   - **Single**
     Choose this option when you want to select a specific source SVM associated with one or more destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, choose this option to associate a specific SVM source to a specific SVM destination that is assigned only to that client.
   - **None (External)**
     Choose this option when you want to create an association between a source SVM and an external flexible volume of a destination SVM.
4. Select one or both of the protection relationship types you want to create:
   - **SnapMirror**
   - **SnapVault**
5. Click **Next**.
6. Select one or more SVM protection destination.
7. Click **Finish**.

Related tasks

Adding a user on page 379

Viewing SVM associations

You can use the Storage Virtual Machine Associations page to view existing Storage Virtual Machine (SVM) associations and their properties and to determine if additional SVM associations are required.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Step

1. Click **Storage > Storage Virtual Machine Associations**.

   The list of SVM associations and their properties is displayed.

Related tasks

Adding a user on page 379

Deleting SVM associations

You can delete Storage Virtual Machine (SVM) associations for partner applications to remove the secondary provisioning relationship between source and destination SVMs; for example, you might do this when the destination SVM is full and you want to create a new SVM protection association.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The **Delete** button is disabled until at least one SVM association is selected. Changing associations using the delete and create functions affects only future provisioning operations; it does not move existing destination volumes.

Steps

1. Click **Storage > Storage Virtual Machine Associations**.
2. Select at least one SVM association.

   The **Delete** button is enabled.
3. Click **Delete.**

   A warning dialog box is displayed.
4. Click **Yes** to continue.

   The selected SVM association is removed from the list.

Related tasks

Adding a user on page 379

What jobs are

A job is a series of tasks that you can monitor using Unified Manager. Viewing jobs and their associated tasks enables you to determine if a they have completed successfully.

Jobs are initiated when you create SnapMirror and SnapVault relationships, when you perform any relationship operation (break, edit, quiesce, remove, resume, resynchronize, and reverse resync), when you perform data restoration tasks, when you log in to a cluster, and so on.

When you initiate a job, you can use the Jobs page and the Job details page to monitor the job and the progress of the associated job tasks.

Monitoring jobs

You can use the Jobs page to monitor job status and to view job properties such as storage service type, state, submitted time, and completed time to determine whether or not a job has successfully completed.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Jobs**.

   The Jobs page is displayed.

2. View the **Status** column to determine the status of those jobs currently running.
3. Click on a job name to view details about that particular job.

   The Job details page is displayed.

Related tasks

Adding a user on page 379

Viewing job details

After you start a job, you can track its progress from the Job details page and monitor the associated tasks for possible errors.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Jobs**.
2. From the **Jobs** page, click a job name in the **Name** column to display the list of tasks associated with the job.
3. Click on a task to display additional information in the **Task Details** pane and the **Task Messages** pane to the right of the task list.

Related tasks

Adding a user on page 379

Aborting jobs

You can use the Jobs page to abort a job if it is taking too long to finish, is encountering too many errors, or is no longer needed. You can abort a job only if its status and type allow it. You can abort any running job.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Jobs**.
2. From the list of jobs, select one job.
3. Click **Abort**.
4. At the confirmation prompt, click **Yes** to abort the selected job.

Related tasks

Adding a user on page 379

Retrying a failed protection job

After you have taken measures to fix a failed protection job, you can use **Retry** to run the job again. Retrying a job creates a new job using the original job ID.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can retry only one failed job at a time. Selecting more than one job disables the **Retry** button. Only jobs of the type Protection Configuration and Protection Relationship Operation can be retried.

Steps

1. Click **Jobs**.
2. From the list of jobs, select a single failed Protection Configuration or Protection Relationship Operation type job.

   The **Retry** button is enabled.
3. Click **Retry**.

   The job is restarted.

Related references

Job details page on page 330

Description of Protection relationships window and dialog boxes

You can view and manage protection-related details such as resource pools, Storage Virtual Machine (SVM) associations, and protection jobs. You can use the appropriate Setup Options dialog box to configure global threshold values for aggregates, volumes, and relationships.

Resource Pools page

The Resource Pools page displays existing resource pools and their members, and enables you to create, monitor, and manage resource pools for provisioning purposes.

- *Command buttons* on page 316
- *Resource Pools list* on page 316
- *Members list command buttons* on page 317
- *Members list* on page 317

Command buttons

The command buttons enable you to perform the following tasks:

Create

Launches the Create Resource Pool dialog box, which you can use to create resource pools.

Edit

Enables you to edit the name and description of the resource pools that you create.

Delete

Enables you to delete one or more resource pools.

Resource Pools list

The Resource Pools list displays (in tabular format) the properties of existing resource pools.

Resource Pool

Displays the name of the resource pool.

Description

Describes the resource pool.

SnapLock Type

Displays the SnapLock type that is being used by the aggregates in the resource pool. Valid values for SnapLock type are Compliance, Enterprise, and Non-SnapLock. A resource pool can contain aggregates of only one SnapLock type.

Total Capacity

Displays the total capacity (in MB, GB, and so on) of the resource pool.

Used Capacity

Displays the amount of space (in MB, GB, and so on) that is used in the resource pool.

Available Capacity

Displays the amount of space (in MB, GB, and so on) that is available in the resource pool.

Used %

Displays the percentage of space that is used in the resource pool.

Members list command buttons

The Members list command buttons enable you to perform the following tasks:

Add

Enables you to add members to the resource pool.

Delete

Enables you to delete one or more members from the resource pool.

Members list

The Members list displays (in tabular format) the resource pool members and their properties when a resource pool is selected.

Status

Displays the current status of the member aggregate. The status can be Critical (✖), Error (❗), Warning (⚠), or Normal (✓).

Aggregate Name

Displays the name of the member aggregate.

State

Displays the current state of the aggregate, which can be one of the following:

- Offline
  Read or write access is not allowed.
- Online
  Read and write access to the volumes that are hosted on this aggregate is allowed.
- Restricted
  Limited operations (such as parity reconstruction) are allowed, but data access is not allowed.
- Creating
  The aggregate is being created.
- Destroying
  The aggregate is being destroyed.
- Failed
  The aggregate cannot be brought online.
- Frozen
  The aggregate is (temporarily) not serving requests.
- Inconsistent
  The aggregate has been marked corrupted; you should contact technical support.
- Iron Restricted
  Diagnostic tools cannot be run on the aggregate.
- Mounting
  The aggregate is in the process of mounting.
- Partial
  At least one disk was found for the aggregate, but two or more disks are missing.

- Quiescing
  The aggregate is being quiesced.
- Quiesced
  The aggregate is quiesced.
- Reverted
  The revert of an aggregate is completed.
- Unmounted
  The aggregate has been unmounted.
- Unmounting
  The aggregate is being taken offline.
- Unknown
  The aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

By default, this column is hidden.

Cluster

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

Node

Displays the name of the node on which the aggregate resides.

Total Capacity

Displays the total capacity (in MB, GB, and so on) of the aggregate.

Used Capacity

Displays the amount of space (in MB, GB, and so on) that is used in the aggregate.

Available Capacity

Displays the amount of space (in MB, GB, and so on) that is available in the aggregate.

Used %

Displays the percentage of space that is used in the aggregate.

Disk Type

Displays the RAID configuration type, which can be one of the following:

- RAID0: All the RAID groups are of type RAID0.
- RAID4: All the RAID groups are of type RAID4.
- RAID-DP: All the RAID groups are of type RAID-DP.
- RAID-TEC: All the RAID groups are of type RAID-TEC.
- Mixed RAID: The aggregate contains RAID groups of different RAID types (RAID0, RAID4, RAID-DP, and RAID-TEC).

By default, this column is hidden.

Related tasks

Viewing resource pools inventory on page 308
Deleting resource pools on page 310
Creating resource pools on page 307
Editing resource pools on page 308
Adding resource pool members on page 309

Removing aggregates from resource pools on page 309

Create Resource Pool dialog box

You can use the Create Resource Pool dialog box to name and describe a new resource pool and to add aggregates to and delete aggregates from that resource pool.

The text boxes enable you to add the following information to create a resource pool:

Resource Pool Name

Enables you to specify a resource pool name.

Description

Enables you to describe a resource pool.

Members

Displays the members of the resource pool. You can also add and delete members.

Command buttons

The command buttons enable you to perform the following tasks:

Add

Opens the Aggregates dialog box so that you can add aggregates from a specific cluster to the resource pool. You can add aggregates from different clusters, but the same aggregates cannot be added to more than one resource pool. Aggregates belonging to clusters that are running Data ONTAP 8.1 are not displayed in the Aggregates dialog box.

Remove

Enables you to remove selected aggregates from the resource pool.

Create

Creates the resource pool. This button is not enabled until information has been entered in the Resource Pool Name or Description fields.

Cancel

Discards the changes and closes the Create Resource Pool dialog box.

Related concepts

Storage Virtual Machine (SVM) and resource pool requirements to support storage services on page 311

Related tasks

Creating resource pools on page 307
Adding resource pool members on page 309

Edit Resource Pool dialog box

You can use the Edit Resource Pool dialog box to change the name and description of an existing resource pool. For example, if the original name and description is inaccurate or incorrect, you can change them so they are more precise.

Text boxes

The text boxes enable you to change the following information for the selected resource pool:

Resource Pool Name

Enables you to enter a new name.

Description

Enables you to enter a new description.

Command buttons

The command buttons enable you to perform the following tasks:

Save

Saves the changes to the resource pool name and description.

Cancel

Discards the changes and closes the Edit Resource Pool dialog box.

Related tasks

Editing resource pools on page 308

Aggregates dialog box

You can use the Aggregates dialog box to select the aggregates that you want to add to your resource pool.

- *Command buttons* on page 320
- *Aggregates list* on page 320

Command buttons

The command buttons enable you to perform the following tasks:

Add

Adds the selected aggregates to the resource pool. The Add button is not enabled until at least one aggregate is selected.

Cancel

Discards the changes, and closes the Aggregates dialog box.

Aggregates list

The Aggregates list displays (in tabular format) the names and properties of monitored aggregates.

Status

Displays the current status of a volume. The status can be Critical (✖), Error (!), Warning (⚠), or Normal (✓).

You can move the pointer over the status to view more information about the event or events generated for the volume.

Aggregate Name

Displays the name of the aggregate.

State

Displays the current state of the aggregate, which can be one of the following:

- Offline
  Read or write access is not allowed.
- Restricted

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

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

Cluster

Displays the name of the cluster on which the aggregate resides.

Node

Displays the name of the storage controller that contains the aggregate.

Total Capacity

Displays the total data size (in MB, GB, and so on) of the aggregate. By default, this column is hidden.

Committed Capacity

Displays the total space (in MB, GB, and so on) that is committed for all the volumes in the aggregate. By default, this column is hidden.

Used Capacity

Displays the amount of space (in MB, GB, and so on) that is used in the aggregate.

Available Capacity

Displays the amount of space (in MB, GB, and so on) that is available for data in the aggregate. By default, this column is hidden.

Available %

Displays the percentage of space that is available for data in the aggregate. By default, this column is hidden.

Used %

Displays the percentage of space that is used by data in the aggregate.

RAID Type

Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, RAID-TEC, or Mixed RAID.

Related tasks

Adding resource pool members on page 309

Storage Virtual Machine Associations page

The Storage Virtual Machine Associations page enables you to view existing Storage Virtual Machine (SVM) associations between source and destination SVMs and to create new SVM associations for use by partner applications to create SnapMirror and SnapVault relationships.

Command buttons

The command buttons enable you to perform the following tasks:

Create

Opens the Create Storage Virtual Machine Associations wizard.

Delete

Enables you to delete the selected SVM associations.

Storage Virtual Machine (SVM) Associations list

The Storage Virtual Machine Associations list displays in a table the source and destination SVM associations that have been created and the type of protection relationship allowed for each association.

Source Storage Virtual Machine

Displays the name of the source SVM.

Source Cluster

Displays the name of the source cluster.

Destination Storage Virtual Machine

Displays the name of the destination SVM.

Destination Cluster

Displays the name of the destination cluster.

Type

Displays the type of protection relationship. Relationship types are either SnapMirror or SnapVault.

Related tasks

Viewing SVM associations on page 313
Creating Storage Virtual Machine (SVM) associations on page 312
Deleting SVM associations on page 313

Create Storage Virtual Machine Associations wizard

The Create Storage Virtual Machine Associations wizard enables you to associate source and destination Storage Virtual Machines (SVMs) for use in SnapMirror and SnapVault protection relationships.

Select Source SVM

The Select Source Storage Virtual Machine panel enables you to select the source, or primary, SVM in the SVM association.

Any

Enables you to create an association between any SVM source to one or more destination, or secondary, SVM. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVM. For example, you might want applications from several different sources at different locations backed up to one or more destination SVM in one location.

Single

Enables you to associate a specific source SVM with one or more destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, choose this option to associate a specific SVM source to a specific SVM destination that is assigned only to that client.

None (External)

Enables you to create an association between a source SVM and an external flexible volume of a destination SVM.

- Storage Virtual Machine
  Lists the names of the available source SVMs
- Cluster
  Lists the clusters on which each SVM resides

Allow these kinds of relationships

Enables you to select the relationship type for the association:

- SnapMirror
  Specifies a SnapMirror relationship as the association type. Selecting this option enables data replication from the selected sources to the selected destinations.
- SnapVault
  Specifies a SnapVault relationship as the association type. Selecting this option enables backups from the selected primary locations to the selected secondary locations.

Select Protection Destinations

The Select Protection Destinations panel of the Create Storage Virtual Machine Associations wizard enables you to select where to copy or replicate the data. You can create an association on only one destination SVM per cluster.

Command buttons

The command buttons enable you to perform the following tasks:

Next

Advances you to the next page in the wizard.

Back

Returns you to the previous page in the wizard.

Finish

Applies your selections and creates the association.

Cancel

Discards the selections and closes the Create Storage Virtual Machine Associations wizard.

Jobs page

The Jobs page enables you to view the current status and other information about all partner application protection jobs that are currently running, as well as jobs that have completed. You can use this information to see which jobs are still running and whether a job has succeeded or failed.

- *Command buttons* on page 324
- *Jobs list* on page 324

Command buttons

The command buttons enable you to perform the following tasks:

Abort

Aborts the selected job. This option is available only if the selected job is running.

Retry

Restarts a failed job of type Protection Configuration or Protection Relationship Operation. You can retry only one failed job at a time. If more than one failed job is selected, the **Retry** button is disabled. You cannot retry failed storage service jobs.

Refresh

Refreshes the list of jobs and the information associated with them.

Jobs list

The Jobs list displays, in tabular format, a list of the jobs that are in progress. By default, the list displays only the jobs generated within the past week. You can use column sorting and filtering to customize which jobs are displayed.

Status

Displays the current status of a job. The status can be Error (⚠) or Normal (✅).

Job Id

Displays the identification number of the job. By default, this column is hidden.

The job identification number is unique and is assigned by the server when it starts the job. You can search for a particular job by entering the job identification number in the text box provided by the column filter.

Name

Displays the name of the job.

Type

Displays the job type. The job types are as follows:

Cluster Acquisition

A Workflow Automation job is rediscovering a cluster.

Protection Configuration

A protection job is initiating Workflow Automation workflows, such as cron schedules, SnapMirror policy creation, and so on.

Protection Relationship Operation

A protection job is running SnapMirror operations.

Protection Workflow Chain

A Workflow Automation job is executing multiple workflows.

Restore

A restore job is running.

Cleanup

The job is cleaning up storage service member artifacts that are no longer needed for restore purposes.

Conform

The job is checking the configuration of storage service members to ensure that they conform.

Destroy

The job is destroying a storage service.

Import

The job is importing unmanaged storage objects into an existing storage service.

Modify

The job is modifying attributes of an existing storage service.

Subscribe

The job is subscribing members to a storage service.

Unsubscribe

The job is unsubscribing members from a storage service.

Update

A protection update job is running.

WFA Configuration

A Workflow Automation job is pushing cluster credentials and synchronizing database caches.

State

Displays the running state of the job. State options are as follows:

Aborted

The job has been aborted.

Aborting

The job is in the process of aborting.

Completed

The job has finished.

Running

The job is running.

Submitted Time

Displays the time the job was submitted.

Duration

Displays the amount of time the job took to complete. This column is displayed by default.

Completed Time

Displays the time the job finished. By default, this column is hidden.

Related tasks

Monitoring jobs on page 314
Viewing job details on page 314
Aborting jobs on page 315

Advanced Secondary Settings dialog box

You can use the Advanced Secondary Settings dialog box to enable version-flexible replication, multiple copy backup, and space-related settings on a secondary volume.

Space-related settings maximize the amount of data being stored, including the following: deduplication, data compression, autogrow, and space guarantee. You might use the Advanced Secondary Settings dialog box when you want to change enable or disable the current settings.

Enable Version-Flexible Replication

Enables SnapMirror with version-flexible replication. Version-flexible replication enables SnapMirror protection of a source volume even if the destination volume is running under an earlier version of Data ONTAP than that of the source volume, as long as both source and destination are running under Data ONTAP 8.3 or higher.

- Enable Backup
  If version-flexible replication is enabled, also enables multiple Snapshot copies of the SnapMirror source data to be transferred to and retained at the SnapMirror destination.

Enable Deduplication

Enables deduplication on the secondary volume in a SnapVault relationship so that duplicate data blocks are eliminated to achieve space savings. You might use deduplication when space savings are at least 10 percent and when data overwrite rate is not rapid. Deduplication is often used for virtualized environments, file shares, and backup data. This setting is disabled by default. When enabled, this operation is initiated after each transfer.

- Enable Compression
  Enables transparent data compression. You might use compression when space savings are at least 10 percent, when the potential overhead is acceptable, and when there are sufficient system resources for compression to complete during nonpeak hours. In a SnapVault relationship, this setting is disabled by default. Compression is available only when deduplication is selected.
    - Compress Inline
      Enables immediate space savings by compressing data before writing data to disk. You might use inline compression when your system has no more than 50 percent utilization during peak hours, and when the system can accommodate new writes and additional CPU during peak hours. This setting is available only when “Enable Compression” is selected.

Enable Autogrow

Enables you to automatically grow the destination volume when the free space percentage is below the specified threshold, as long as space is available on the associated aggregate.

Maximum Size

Sets the maximum percentage to which a volume can grow. The default is 20 percent greater than the source volume size. A volume does not grow automatically if the current size is greater than or equal to the maximum autogrow percentage. This field is enabled only when the autogrow setting is enabled.

Increment Size

Specifies the percentage increment by which the volume automatically grows before reaching the maximum percentage of the source volume.

Space Guarantee

Ensures that enough space is allocated on the secondary volume so that data transfers always succeed. The space guarantee setting can be one of the following:

- File
- Volume
- None

For example, you might have a 200 GB volume that contains files totaling 50 GB; however, those files hold only 10 GB of data. Volume guarantee allocates 200 GB to the destination volume, regardless of contents on the source. File guarantee allocates 50 GB to ensure that enough space is reserved for files on the source; selecting None in this scenario means that only 10 GB is allocated on the destination for the actual space used by file data on the source.

The space guarantee is set to Volume by default.

Command buttons

The command buttons enable you to perform the following tasks:

Apply

Saves the selected efficiency settings and applies them when you click **Apply** in the Configure Protection dialog box.

Cancel

Discards your selections and closes the Advanced Destination Settings dialog box.

Related references

Configure Protection dialog box on page 337

Related information

NetApp Technical Report 3966: NetApp Data Compression and Deduplication Deployment and Implementation Guide (Clustered Data ONTAP)

Advanced Destination Settings dialog box

You can use the Advanced Destination Settings dialog box to enable space guarantee settings on a destination volume. You might select advanced settings when space guarantee is disabled on the source, but you want it enabled on the destination. The settings for deduplication, compression, and autogrow in a SnapMirror relationship are inherited from the source volume and cannot be changed.

Space Guarantee

Ensures that enough space is allocated on the destination volume so that data transfers always succeed. The space guarantee setting can be one of the following:

- File

Space guarantee for files is not available in Data ONTAP 8.3.

- Volume
- None

For example, you might have a 200-GB volume that contains files totaling 50 GB; however, those files hold only 10 GB of data. Volume guarantee allocates 200 GB to the destination volume, regardless of contents on the source. File guarantee allocates 50 GB to ensure that enough space is reserved for source files on the destination; selecting **None** in this scenario means that only 10 GB is allocated on the destination for the actual space used by file data on the source.

The space guarantee is set to Volume by default.

Related references

Configure Protection dialog box on page 337

Restore dialog box

You can use the Restore dialog box to restore data to a volume from a specific Snapshot copy. Restoring data on volumes using a Data ONTAP version earlier than 8.2 is not supported.

Restore from

The Restore from area enables you to specify from where you want to restore data.

Volume

Specifies the volume from which you want to restore data. By default, the volume on which you initiated the restore action is selected, or you can select a different volume from the drop-down list that contains all the volumes with protection relationships to the volume on which you initiated the restore action.

Snapshot copy

Specifies which Snapshot copy you want to use to restore data. By default, the most recent Snapshot copy is selected. You can also select a different Snapshot copy from the drop-down list. The Snapshot copy list changes according to which volume is selected.

Select items to restore

The Select items to restore area enables you to select either the entire volume or specific files and folders you want to restore. You can select a maximum of 10 files, folders, or a combination of both. When the maximum number of items is selected, the item selection check boxes are disabled.

Path field

Displays the path to the data you want to restore. You can either navigate to the folder and files you want to restore, or you can type the path. This field is empty until you select or type a path. Clicking ↰ after you have chosen a path moves you up one level in the directory structure.

Folders and files list

Displays the contents of the path you entered. By default, the root folder is initially displayed. Clicking a folder name displays the contents of the folder.

You can select items to restore as follows:

- When you enter the path with a particular file name specified in the path field, the specified file is displayed in the Folders and Files.
- When you enter a path without specifying a particular file, the contents of the folder are displayed in the Folders and Files list, and you can select up to 10 files, folders, or a combination of both to

restore. However, if a folder contains more than 995 items, a message displays to indicate there are too many items to display and if you proceed with the operation, all items in the specified folder are restored.

Note: You cannot restore NTFS file streams.

Restore to

The Restore to area enables you to specify where you want to restore the data.

Original Location in `Volume_Name`

Restores the selected data to the directory on the source from which the data was originally backed up.

Alternate Location

Restores the selected data to a new location:

- Restore Path

  Specifies an alternate path for restoring the selected data. The path must already exist. You can use the **Browse** button to navigate to the location where you want the data restored, or you can enter the path manually using the format cluster://svm/volume/path.

- Preserve directory hierarchy

  When checked, preserves the structure of the original file or directory. For example, if the source is /A/B/C/myFile.txt and the destination is /X/Y/Z, Unified Manager restores the data using the following directory structure on the destination: /X/Y/Z/A/B/C/myFile.txt.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards your selections and closes the Restore dialog box.

Restore

Applies your selections and begins the restore process.

Related tasks

Restoring data using the Volumes page on page 305

Restoring data using the Volume details page on page 306

Related references

Browse Directories dialog box on page 329

Browse Directories dialog box

You can use the Browse Directories dialog box when you want to restore data to a directory on a cluster and Storage Virtual Machine (SVM) that is different from the original source. The original source cluster and volume are selected by default.

The Browse Directories dialog box enables you to select the cluster, SVM, volume, and directory path to which you want data restored.

Cluster

Lists the available cluster destinations to which you can restore. By default, the cluster of the original source volume is selected.

SVM drop-down list

Lists the available SVMs available for the selected cluster. By default, the SVM of the original source volume is selected.

Volume

Lists all of the read/write volumes in a selected SVM. You can filter the volumes by name and by space available. The volume with the most space is listed first, and so on, in descending order. By default, the original source volume is selected.

File path text box

Enables you to type the file path to which you want data restored. The path you enter must already exist.

Name

Displays the names of the available folders for the selected volume. Clicking a folder in the Name list displays the subfolders, if any exist. Files contained in the folders are not displayed. Clicking ↰ after you have selected a folder moves you up one level in the directory structure.

Command buttons

The command buttons enable you to perform the following tasks:

Select Directory

Applies your selections and closes the Browse Directories dialog box. If no directory is selected, this button is disabled.

Cancel

Discards your selections and closes the Browse Directories dialog box.

Related references

Restore dialog box on page 328

Job details page

The Job details page enables you to view status and other information about specific protection job tasks that are running, that are queued, or that have completed. You can use this information to monitor protection job progress and to troubleshoot job failures.

Job summary

The job summary displays the following information:

- Job ID
- Type
- State
- Submitted Time
- Completed Time
- Duration

Command buttons

The command buttons enable you to perform the following tasks:

Refresh

Refreshes the task list and the properties associated with each task.

View Jobs

Returns you to the Jobs page.

Job tasks list

The Job tasks list displays in a table all the tasks associated with a specific job and the properties related to each task.

Started Time

Displays the day and time the task started. By default, the most recent tasks are displayed at the top of the column and older tasks are displayed at the bottom.

Type

Displays the type of task.

State

The state of a particular task:

Completed

The task has finished.

Queued

The task is about to run.

Running

The task is running.

Waiting

A job has been submitted and some associated tasks are waiting to be queued and executed.

Status

Displays the task status:

Error (⚠)

The task failed.

Normal (✓)

The task succeeded.

Skipped (↻)

A task failed, resulting in subsequent tasks being skipped.

Duration

Displays the elapsed time since the task began.

Completed Time

Displays the time the task completed. By default, this column is hidden.

Task ID

Displays the GUID that identifies an individual task for a job. The column can be sorted and filtered. By default, this column is hidden.

Dependency order

Displays an integer representing the sequence of tasks in a graph, with zero being assigned to the first task. By default, this column is hidden.

Task Details pane

Displays additional information about each job task, including the task name, task description, and, if the task failed, a reason for the failure.

Task Messages pane

Displays messages specific to the selected task. Messages might include a reason for the error and suggestions for resolving it. Not all tasks display task messages.

Related tasks

Viewing job details on page 314

Lag Thresholds for Unmanaged Relationships dialog box

The Lag Thresholds for Unmanaged Relationships dialog box enables you to configure global lag warning and error threshold values for unmanaged protection relationships so that you are notified and can take action when lag or threshold errors occur. Changes to these settings are applied during the next scheduled update.

You must have the OnCommand Administrator or Storage Administrator role.

Events are generated when a threshold is breached. You can take corrective actions for such events. Lag threshold settings for unmanaged relationships are enabled by default.

- *Lag Thresholds for Unmanaged Relationships* on page 332
- *Command buttons* on page 332

Lag Thresholds for Unmanaged Relationships area

The Lag area enables you set unmanaged relationship lag thresholds for the following conditions:

Warning

Specifies the percentage at which the lag duration equals or exceeds the lag warning threshold:

- Default value: 150 percent
- Events generated: SnapMirror Relationship Lag Warning or SnapVault Relationship Lag Warning
- Event severity: Warning

Error

Specifies the percentage at which the lag duration equals or exceeds the lag error threshold:

- Default value: 250 percent
- Events generated: SnapMirror Relationship Lag Error or SnapVault Relationship Lag Error
- Event severity: Error

Command buttons

The command buttons enable you to save or cancel the setup options:

Restore to Factory Defaults

Enables you to restore the configuration settings to the factory default values.

Save

Saves the configuration settings for the selected option.

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Editing lag threshold settings for unmanaged protection relationships on page 127

Edit Aggregate Thresholds dialog box

You can configure alerts to send notifications when an event related to an aggregate's capacity is generated, and you can take corrective actions for the event. For example, for the Aggregate Full threshold, you can configure an alert to send notification when the condition persists over a specified period.

You must have the OnCommand Administrator or Storage Administrator role.

The Edit Aggregate Thresholds dialog box enables you to configure aggregate-level thresholds that are applied to selected aggregates. If you configure aggregate-level thresholds, they take priority over the global-level threshold values. You can configure threshold settings for capacity, growth, and Snapshot copies at the aggregate level. If these settings are not configured, the global threshold values are applied.

Note: The threshold values are not applicable to the root aggregate of the node.

- *Capacity* on page 333
- *Growth* on page 334
- *Snapshot Copies* on page 334
- *Command buttons* on page 334

Capacity area

The Capacity area enables you to set the following aggregate capacity threshold conditions:

Space Nearly Full

Specifies the percentage at which an aggregate is considered to be nearly full. It also displays the size of the aggregate corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which an aggregate is considered full. It also displays the size of the aggregate corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Nearly Overcommitted

Specifies the percentage at which an aggregate is considered to be nearly overcommitted.

Overcommitted

Specifies the percentage at which an aggregate is considered to be overcommitted.

Days Until Full

Specifies the number of days remaining before the aggregate reaches full capacity.

Growth

The Growth area enables you to set the following threshold condition for aggregate growth:

Growth Rate

Specifies the percentage at which an aggregate's growth rate is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event.

Growth Rate Sensitivity

Specifies the factor that is applied to the standard deviation of an aggregate's growth rate. If the growth rate exceeds the factored standard deviation, an Aggregate Growth Rate Abnormal event is generated.

A lower value for growth rate sensitivity indicates that the aggregate is highly sensitive to changes in the growth rate.

Note: If you modify the growth rate sensitivity for aggregates at the global threshold level, the change is also applied to the growth rate sensitivity for volumes at the global threshold level.

Snapshot copies area

The Snapshot copies area enables you to set the following Snapshot reserve threshold conditions:

Snapshot Reserve Full

Specifies the percentage at which an aggregate has consumed all its space reserved for Snapshot copies.

You can also use the slider to set the threshold value.

Command buttons

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

Restore to Defaults

Enables you to restore the aggregate-level threshold values to the global values.

Save

Saves all the threshold settings.

Save and Close

Saves all the threshold settings and then closes the Edit Aggregate Thresholds dialog box.

Cancel

Ignores the changes (if any) to the threshold settings and closes the Edit Aggregate Thresholds dialog box.

Related concepts

Understanding capacity events and thresholds for node root aggregates on page 145

Related tasks

Editing aggregate threshold settings on page 128

Edit Volume Thresholds dialog box

You can configure alerts to send notifications when an event related to a volume's capacity is generated, and you can take corrective actions for the event. For example, for the Volume Full threshold, you can configure an alert to send notification when the condition persists over a specified period.

You must have the OnCommand Administrator or Storage Administrator role.

The Edit Volume Thresholds dialog box enables you to configure volume-level thresholds that are applied to the selected volumes. When thresholds are configured at the volume level, they take priority over the group-level thresholds or the global-level threshold values.

You can configure threshold settings for capacity, Snapshot copies, qtree quota, growth, and inodes at the volume level. When a group action of volume threshold type is configured, the group action threshold values are used for settings that are not configured at the volume level. When no group action of volume threshold type is configured, areas in Edit Volume Thresholds dialog box that are not configured, use global threshold values.

- *Capacity* on page 335
- *Snapshot Copies* on page 335
- *Quota* on page 336
- *Growth* on page 336
- *Inodes* on page 336
- *Command buttons* on page 336

Capacity area

The Capacity area enables you to set the following volume capacity threshold conditions:

Space Nearly Full

Specifies the percentage at which a volume is considered to be nearly full. It also displays the size of the volume corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which a volume is considered full. It also displays the size of the volume corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Days Until Full

Specifies the number of days remaining before the volume reaches full capacity.

Snapshot Copies

The Snapshot Copies area enables you to set the following threshold conditions for the Snapshot copies in the volume.

Snapshot Reserve Full

Specifies the percentage at which the space reserved for Snapshot copies is considered full.

Days Until Full

Specifies the number of days remaining before the space reserved for Snapshot copies reaches full capacity.

Count

Specifies the number of Snapshot copies that can be created on a volume before the system generates the Too Many Snapshot Copies event.

Note: This field is applicable only for volumes in a cluster running Data ONTAP 8.2 or later.

Quota area

The Quota area enables you to set the following volume quota threshold conditions:

Nearly Overcommitted

Specifies the percentage at which a volume is considered to be nearly overcommitted by qtree quotas.

Overcommitted

Specifies the percentage at which a volume is considered to be overcommitted by qtree quotas.

Growth area

The Growth area enables you to set the following threshold condition for volume growth:

Growth Rate

Specifies the percentage at which a volume's growth rate is considered to be normal before the system generates a Volume Growth Rate Abnormal event.

Growth Rate Sensitivity

Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated.

A lower value for growth rate sensitivity indicates that the volume is highly sensitive to changes in the growth rate.

Note: If you modify the growth rate sensitivity for volumes at the global threshold level, the change is also applied to the growth rate sensitivity for aggregates at the global threshold level.

Inodes

The Inodes enables you to set the following threshold conditions for inodes:

Nearly Full

Specifies the percentage at which a volume is considered to have consumed most of its inodes.

You can also use the sliders to set the threshold value.

Full

Specifies the percentage at which a volume is considered to have consumed all of its inodes.

You can also use the sliders to set the threshold value.

Command buttons

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

Restore to Defaults

Enables you to restore the threshold values to one of the following:

- Group values, if the volume belongs to a group and that group has a volume threshold action type.
- Global values, if the volume does not belong to any group or if it belongs to a group that does not have a volume threshold action type.

Save

Saves all the threshold settings.

Save and Close

Saves all the threshold settings and then closes the Aggregates dialog box.

Cancel

Ignores the changes (if any) to the threshold settings and closes the Aggregates dialog box.

Related tasks

Editing volume threshold settings on page 129

Configure Protection dialog box

You can use the Configure Protection dialog box to create SnapMirror and SnapVault relationships for all read, write, and data protection volumes on clusters that are running Data ONTAP 8.2 or later to ensure that the data on a source volume or primary volume is replicated.

- *Source tab* on page 337
- *SnapMirror tab* on page 338
- *SnapVault tab* on page 339
- *Command buttons* on page 341

Source tab

Topology view

Displays a visual representation of the relationship that you are creating. The source in the topology is highlighted by default.

Source Information

Displays details about the selected source volumes, including the following information:

- Source cluster name
- Source SVM name
- Cumulative volume total size
  Displays the total size of all the source volumes that are selected.
- Cumulative volume used size
  Displays the cumulative volume used size for all the selected source volumes.
- Source volume
  Displays the following information in a table :
  - Source Volume
    Displays the names of the selected source volumes.
  - Type
    Displays the volume type.
  - SnapLock Type
    Displays the SnapLock type of the volume. The options are Compliance, Enterprise, and Non-SnapLock.
  - Snapshot Copy
    Displays the Snapshot copy that is used for the baseline transfer. If the source volume is read/write, the value of Default in the Snapshot copy column indicates that a new Snapshot copy is created by default, and is used for the baseline transfer. If the source volume is a data protection volume, the value of Default in the

Snapshot copy column indicates that no new Snapshot copy is created, and all existing Snapshot copies are transferred to the destination. Clicking the Snapshot copy value displays a list of Snapshot copies from which you can select an existing Snapshot copy to use for the baseline transfer. You cannot select a different default Snapshot copy if the source type is data protection.

SnapMirror tab

Enables you to specify a destination cluster, Storage Virtual Machine (SVM), and aggregate for a protection relationship, as well as a naming convention for destinations while creating a SnapMirror relationship. You can also specify a SnapMirror policy and schedule.

Topology view

Displays a visual representation of the relationship that you are creating. The SnapMirror destination resource in the topology is highlighted by default.

Destination Information

Enables you to select the destination resources for a protection relationship:

- Advanced link
  Launches the Advanced Destination Settings dialog box when you are creating a SnapMirror relationship.
- Cluster
  Lists the clusters that are available as protection destination hosts. This field is required.
- Storage Virtual Machine (SVM)
  Lists the SVMs that are available on the selected cluster. A cluster must be selected before the SVM list is populated. This field is required.
- Aggregate
  Lists the aggregates that are available on the selected SVM. A cluster must be selected before the Aggregate list is populated. This field is required. The Aggregate list displays the following information:
  - Rank
    When multiple aggregates satisfy all the requirements for a destination, the rank indicates the priority in which the aggregate is listed, according to the following conditions:
    1. An aggregate that is located on a different node than the source volume node is preferred to enable fault domain separation.
    2. An aggregate on a node with fewer volumes is preferred to enable load balancing across nodes in a cluster.
    3. An aggregate that has more free space than other aggregates is preferred to enable capacity balancing.

    A rank of 1 means that the aggregate is the most preferred according to the three criteria.
  - Aggregate Name
    Name of the aggregate
  - Available Capacity
  - Amount of space that is available on the aggregate for data
  - Resource Pool
    Name of the resource pool to which the aggregate belongs

- Naming Convention

  Specifies the default naming convention that is applied to the destination volume. You can accept the naming convention that is provided, or you can create a custom one. The naming convention can have the following attributes: %C, %M, %V, and %N, where %C is the cluster name, %M is the SVM name, %V is the source volume, and %N is the topology destination node name.

  The naming convention field is highlighted in red if your entry is invalid. Clicking the “Preview Name” link displays a preview of the naming convention that you entered, and the preview text updates dynamically as you type a naming convention in the text field. A suffix between 001 and 999 is appended to the destination name when the relationship is created, replacing the *nnn* that displays in the preview text, with 001 being assigned first, 002 assigned second, and so on.

Relationship Settings

Enables you to specify the maximum transfer rate, SnapMirror policy, and schedule that the protection relationship uses:

- Max Transfer Rate

  Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running Data ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored.

- SnapMirror Policy

  Specifies the ONTAP SnapMirror policy for the relationship. The default is DPDefault.

- Create Policy

  Launches the Create SnapMirror Policy dialog box, which enables you to create and use a new SnapMirror policy.

- SnapMirror Schedule

  Specifies the ONTAP SnapMirror policy for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.

- Create Schedule

  Launches the Create Schedule dialog box, which enables you to create a new SnapMirror schedule.

SnapVault tab

Enables you to specify a secondary cluster, SVM, and aggregate for a protection relationship, as well as a naming convention for secondary volumes while creating a SnapVault relationship. You can also specify a SnapVault policy and schedule.

Topology view

Displays a visual representation of the relationship that you are creating. The SnapVault secondary resource in the topology is highlighted by default.

Secondary Information

Enables you to select the secondary resources for a protection relationship:

- Advanced link

  Launches the Advanced Secondary Settings dialog box.

- Cluster

  Lists the clusters that are available as secondary protection hosts. This field is required.

- Storage Virtual Machine (SVM)
  Lists the SVMs that are available on the selected cluster. A cluster must be selected before the SVM list is populated. This field is required.

- Aggregate
  Lists the aggregates that are available on the selected SVM. A cluster must be selected before the Aggregate list is populated. This field is required. The Aggregate list displays the following information:

  - Rank
    When multiple aggregates satisfy all the requirements for a destination, the rank indicates the priority in which the aggregate is listed, according to the following conditions:

    1. An aggregate that is located on a different node than the primary volume node is preferred to enable fault domain separation.
    2. An aggregate on a node with fewer volumes is preferred to enable load balancing across nodes in a cluster.
    3. An aggregate that has more free space than other aggregates is preferred to enable capacity balancing.

    A rank of 1 means that the aggregate is the most preferred according to the three criteria.

  - Aggregate Name
    Name of the aggregate
  - Available Capacity
  - Amount of space that is available on the aggregate for data
  - Resource Pool
    Name of the resource pool to which the aggregate belongs

- Naming Convention
  Specifies the default naming convention that is applied to the secondary volume. You can accept the naming convention that is provided, or you can create a custom one. The naming convention can have the following attributes: %C, %M, %V, and %N, where %C is the cluster name, %M is the SVM name, %V is the source volume, and %N is the topology secondary node name.
  The naming convention field is highlighted in red if your entry is invalid. Clicking the “Preview Name” link displays a preview of the naming convention that you entered, and the preview text updates dynamically as you type a naming convention in the text field. If you type an invalid value, the invalid information displays as red question marks in the preview area. A suffix between 001 and 999 is appended to the secondary name when the relationship is created, replacing the *nnn* that displays in the preview text, with 001 being assigned first, 002 assigned second, and so on.

Relationship Settings

Enables you to specify the maximum transfer rate, SnapVault policy, and SnapVault schedule that the protection relationship uses:

- Max Transfer Rate
  Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running Data ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored.

- SnapVault Policy
  Specifies the ONTAP SnapVault policy for the relationship. The default is XDPDefault.
- Create Policy
  Launches the Create SnapVault Policy dialog box, which enables you to create and use a new SnapVault policy.
- SnapVault Schedule
  Specifies the ONTAP SnapVault schedule for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.
- Create Schedule
  Launches the Create Schedule dialog box, which enables you to create a SnapVault schedule.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards your selections, and closes the Configure Protection dialog box.

Apply

Applies your selections, and begins the protection process.

Related tasks

Creating a SnapVault protection relationship from the Volumes page on page 281
Creating a SnapVault protection relationship from the Volume details page on page 282
Creating a SnapMirror protection relationship from the Volumes page on page 285
Creating a SnapMirror protection relationship from the Volume details page on page 283
Creating cascade or fanout relationships to extend protection from an existing protection relationship on page 290

Related references

Advanced Secondary Settings dialog box on page 326
Advanced Destination Settings dialog box on page 327

Create Schedule dialog box

The Create Schedule dialog box enables you to create a basic or advanced protection schedule for SnapMirror and SnapVault relationship transfers. You might create a new schedule to increase the frequency of data transfers due to frequent data updates, or you might create a less frequent schedule when data changes infrequently.

Destination Cluster

The name of the cluster you selected in the SnapVault tab or SnapMirror tab of the Configure Protection dialog box.

Schedule Name

The name you provide for the schedule. Schedule names can consist of the characters A through Z, a through z, 0 through 9, as well as any of the following special characters: ! @ # $ % ^ & * ( ) _ -. Schedule names may not include the following characters: < >.

Basic or Advanced

The schedule mode you want to use.

Basic mode includes the following elements:

- Repeat
  How often a scheduled transfer occurs. Choices include hourly, daily, and weekly.
- Day
  When a repeat of weekly is selected, the day of the week a transfer occurs.
- Time
  When Daily or Weekly is selected, the time of day a transfer occurs.

Advanced mode includes the following elements:

- Months
  A comma-separated numerical list representing the months of the year. Valid values are 0 through 11, with zero representing January, and so on. This element is optional. Leaving the field blank implies that transfers occur every month.
- Days
  A comma-separated numerical list representing the day of the month. Valid values are 1 through 31. This element is optional. Leaving the field blank implies that a transfer occurs every day of the month.
- Weekdays
  A comma-separated numerical list representing the days of the week. Valid values are 0 through 6, with 0 representing Sunday, and so on. This element is optional. Leaving the field blank implies that a transfer occurs every day of the week. If a day of the week is specified but a day of the month is not specified, a transfer occurs only on the specified day of the week and not every day.
- Hours
  A comma-separated numerical list representing the number of hours in a day. Valid values are 0 through 23, with 0 representing midnight. This element is optional.
- Minutes
  A comma-separated numerical list representing the minutes in an hour. Valid values are 0 through 59. This element is required.

Related tasks

Creating SnapMirror and SnapVault schedules on page 289

Create SnapMirror Policy dialog box

The Create SnapMirror Policy dialog box enables you to create a policy to set the priority for SnapMirror transfers. You use policies to maximize the efficiency of transfers from the source to the destination.

Destination Cluster

The name of the cluster you selected in the SnapMirror tab of the Configure Protection dialog box.

Destination Storage Virtual Machine (SVM)

The name of the SVM you selected in the SnapMirror tab of the Configure Protection dialog box.

Policy Name

The name you provide for the new policy. Policy names can consist of the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underscore (_).

Transfer Priority

The priority at which the transfer is run. You can select either Normal or Low. Transfer relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority.

Comment

An optional field in which you can add comments about the policy.

Transfer Restart

Indicates what restart action to take when a transfer is interrupted by an abort operation or any type of failure, such as a network outage. You can select one of the following:

- Always
  Specifies that a new Snapshot copy is created before restarting a transfer, then, if one exists, the transfer is restarted from a checkpoint, followed by an incremental transfer from the newly created Snapshot copy..
- Never
  Specifies that interrupted transfers are never restarted.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards the selections and closes the Configure Protection dialog box.

Apply

Applies your selections and begins the protection process.

Related tasks

Creating SnapMirror policies on page 292

Create SnapVault Policy dialog box

The Create SnapVault Policy dialog box enables you to create a policy to set the priority for SnapVault transfers. You use policies to maximize the efficiency of transfers from the primary to the secondary volume.

Destination Cluster

The name of the cluster that you selected in the SnapVault tab of the Configure Protection dialog box.

Destination Storage Virtual Machine (SVM)

The name of the SVM that you selected in the SnapVault tab of the Configure Protection dialog box.

Policy Name

The name you provide for the new policy. Policy names can consist of the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underscore (_).

Transfer Priority

The priority at which the transfer is run. You can select either Normal or Low. Transfer relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority. The default setting is Normal.

Comment

An optional field in which you can a add comment of up to 255 characters about the SnapVault policy.

Ignore Access Time

Specifies whether incremental transfers are ignored for files that have only their access time changed.

Replication Label

Lists in a table the rules associated with Snapshot copies selected by Data ONTAP that have a specific replication label in a policy. The following information and actions are also available:

- Command buttons
  The command buttons enable you to perform the following actions:
  - Add
    Enables you to create a Snapshot copy label and retention count.
  - Edit Retention Count
    Enables you to change the retention count for an existing Snapshot copy label. The retention count must be a number between 1 and 251. The sum of all retention counts for all rules cannot exceed 251.
  - Delete
    Enables you to delete an existing Snapshot copy label.
- Snapshot Copy Label
  Displays the Snapshot copy label. If you select one or more volumes with the same local Snapshot copy policy, an entry for each label in the policy is displayed. If you select multiple volumes that have two or more local Snapshot copy policies, the table displays all labels from all policies
- Schedule
  Displays the schedule associated with each Snapshot copy label. If a label has more than one schedule associated with it, the schedules for that label are displayed in a comma-separated list. If you select multiple volumes with the same label but with different schedules, the schedule displays “Various” to indicate that more than one schedule is associated with the selected volumes.
- Destination Retention Count
  Displays the number of Snapshot copies with the specified label that are retained on the SnapVault secondary. Retention counts for labels with multiple schedules displays the sum of retention counts of each label and schedule pair. If you select multiple volumes with two or more local Snapshot copy policies, the retention count is empty.

Related tasks

Creating SnapVault policies on page 293

Edit Relationship dialog box

You can edit an existing protection relationship to change the maximum transfer rate, the protection policy, or the protection schedule.

Destination Information

Destination Cluster

The name of the selected destination cluster.

Destination SVM

The name of the selected Storage Virtual Machine (SVM)

Relationship Settings

Enables you to specify the maximum transfer rate, SnapMirror policy, and schedule that the protection relationship uses:

- Max Transfer Rate
  Specifies the maximum rate at which baseline data is transferred between clusters over the network. When selected, network bandwidth is limited to the value you specify. You can enter a numerical value and then select either kilobytes per second (KBps), megabytes per second (MBps), gigabytes per second (GBps), or terabytes per second (TBps). The maximum transfer rate that you specify must be greater than 1 KBps and less than 4 TBps. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. If the primary cluster and the secondary cluster are the same, this setting is disabled.
- SnapMirror Policy
  Specifies the Data ONTAP SnapMirror policy for the relationship. The default is DPDefault.
- Create Policy
  Launches the Create SnapMirror Policy dialog box, which enables you to create and use a new SnapMirror policy.
- SnapMirror Schedule
  Specifies the Data ONTAP SnapMirror policy for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.
- Create Schedule
  Launches the Create Schedule dialog box, which enables you to create a new SnapMirror schedule.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards the selections and closes the Configure Protection dialog box.

Submit

Applies your selections and closes the Edit Relationship dialog box.

Related tasks

Editing protection relationships on page 291
Editing protection relationships from the Volume details page on page 291

Initialize/Update dialog box

The Initialize/Update dialog box enables you to perform a first-time baseline transfer on a new protection relationship, or to update a relationship if it is already initialized and you want to perform a manual, unscheduled, incremental update.

Transfer Options tab

The Transfer Options tab enables you to change the initialization priority of a transfer and to change the bandwidth used during transfers.

Transfer Priority

The priority at which the transfer is run. You can select either Normal or Low. Relationships with policies that specify a normal transfer priority run before those that specify a low transfer priority. Normal is selected by default.

Max Transfer Rate

Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running Data ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored. If you select more than one relationship with different maximum transfer rates, you can specify one of the following maximum transfer rate settings:

- Use values specified during individual relationship setup or edit

  When selected, initialization and update operations use the maximum transfer rate specified at the time of each relationship's creation or edit. This field is available only when multiple relationships with different transfer rates are being initialized or updated.
- Unlimited

  Indicates that there is no bandwidth limitation on transfers between relationships. This field is available only when multiple relationships with different transfer rates are being initialized or updated.
- Limit bandwidth to

  When selected, network bandwidth is limited to the value you specify. You can enter a numerical value and then select either kilobytes per second (KBps), Megabytes per second (MBps), Gigabytes per second (GBps), or Terabytes per second (TBps). The maximum transfer rate that you specify must be greater than 1 KBps and less than 4 TBps.

Source Snapshot Copies tab

The Source Snapshot Copies tab displays the following information about the source Snapshot copy that is used for the baseline transfer:

Source Volume

Displays the names of the corresponding source volumes.

Destination Volume

Displays the names of the selected destination volumes.

Source Type

Displays the volume type. The type can be either Read/write or Data Protection.

Snapshot Copy

Displays the Snapshot copy that is used for the data transfer. Clicking the Snapshot copy value displays the Select Source Snapshot Copy dialog box, in which you can select a specific Snapshot copy for your transfer, depending on the type of protection relationship that you have and the operation that you are performing. The option to specify a different Snapshot copy is not available for data protection type sources.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards your selections and closes the Initialize/Update dialog box.

Submit

Saves your selections and starts the initialize or update job.

Related tasks

Initializing or updating protection relationships from the Volume details page on page 299
Initializing or updating protection relationships on page 300

Related references

Select Source Snapshot Copy dialog box on page 348

Resynchronize dialog box

The Resynchronize dialog box enables you to resynchronize data on a SnapMirror or SnapVault relationship that was previously broken and then the destination was made a read/write volume. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Resynchronization Options tab

The Resynchronization Options tab enables you to set the transfer priority and the maximum transfer rate for the protection relationship that you are resynchronizing.

Transfer Priority

The priority at which the transfer is run. You can select either Normal or Low. Relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority.

Max Transfer Rate

Specifies the maximum rate at which data is transferred between clusters over the network. When selected, network bandwidth is limited to the value that you specify. You can enter a numerical value and then select either kilobytes per second (KBps), megabytes per second (MBps), gigabytes per second (GBps), or TBps. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running Data ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is disabled.

Source Snapshot Copies tab

The Source Snapshot Copies tab displays the following information about the source Snapshot copy that is used for the baseline transfer:

Source Volume

Displays the names of the corresponding source volumes.

Destination Volume

Displays the names of the selected destination volumes.

Source Type

Displays the volume type: either Read/write or Data Protection.

Snapshot Copy

Displays the Snapshot copy that is used for the data transfer. Clicking the Snapshot copy value displays the Select Source Snapshot Copy dialog box, in which can select a specific Snapshot copy for your transfer, depending on the type of protection relationship you have and the operation you are performing.

Command buttons

Submit

Begins the resynchronization process and closes the Resynchronize dialog box.

Cancel

Cancels your selections and closes the Resynchronize dialog box.

Related tasks

Resynchronizing protection relationships from the Volume details page on page 301
Resynchronizing protection relationships on page 302

Related references

Select Source Snapshot Copy dialog box on page 348

Select Source Snapshot Copy dialog box

You use the Select Source Snapshot Copy dialog box to select a specific Snapshot copy to transfer data between protection relationships, or you select the default behavior, which varies depending on whether you are initializing, updating, or resynchronizing a relationship, and whether the relationship is a SnapMirror or SnapVault.

Default

Enables you to select the default behavior for determining which Snapshot copy is used for initialize, update, and resynchronize transfers for SnapVault and SnapMirror relationships.

If you are performing a SnapVault transfer, the default behavior for each operation is as follows:

Operation	Default SnapVault behavior when source is read/write	Default SnapVault behavior when source is Data Protection (DP)
Initialize	Creates a new Snapshot copy and transfers it.	Transfers the last exported Snapshot copy.
Update	Transfers only labeled Snapshot copies, as specified in the policy.	Transfers the last exported Snapshot copy.
Resynchronize	Transfers all labeled Snapshot copies created after the newest common Snapshot copy.	Transfers the newest labeled Snapshot copy.

If you are performing a SnapMirror transfer, the default behavior for each operation is as follows:

Operation	Default SnapMirror behavior	Default SnapMirror behavior when relationship is second hop in a SnapMirror to SnapMirror cascade
Initialize	Creates a new Snapshot copy and transfers it and all Snapshot copies created prior to the new Snapshot copy.	Transfers all Snapshot copies from the source.
Update	Creates a new Snapshot copy and transfers it and all Snapshot copies created prior to the new Snapshot copy.	Transfers all Snapshot copies.

Operation	Default SnapMirror behavior	Default SnapMirror behavior when relationship is second hop in a SnapMirror to SnapMirror cascade
Resynchronize	Creates a new Snapshot copy and then transfers all Snapshot copies from the source.	Transfers all Snapshot copies from the secondary volume to the tertiary volume, and deletes any data added after creation of the newest common Snapshot copy.

Existing Snapshot Copy

Enables you to select an existing Snapshot copy from the list if Snapshot copy selection is allowed for that operation.

Snapshot Copy

Displays the existing Snapshot copies from which you can select for a transfer.

Date Created

Displays the date and time the Snapshot copy was created. Snapshot copies are listed from most recent to least recent, with the most recent at the top of the list.

If you are performing a SnapVault transfer and you want to select an existing Snapshot copy to transfer from a source to a destination, the behavior for each operation is as follows:

Operation	SnapVault behavior when specifying a Snapshot copy	SnapVault behavior when specifying a Snapshot copy in a cascade
Initialize	Transfers the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Update	Transfers the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Resynchronize	Transfers the selected Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

If you are performing a SnapMirror transfer and you want to select an existing Snapshot copy to transfer from a source to a destination, the behavior for each operation is as follows:

Operation	SnapMirror behavior when specifying a Snapshot copy	SnapMirror behavior when specifying a Snapshot copy in a cascade
Initialize	Transfers all Snapshot copies on the source, up to the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Update	Transfers all Snapshot copies on the source, up to the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

Operation	SnapMirror behavior when specifying a Snapshot copy	SnapMirror behavior when specifying a Snapshot copy in a cascade
Resynchronize	Transfers all Snapshot copies from the source, up to the selected Snapshot copy, and then deletes any data added after creation of the newest common Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

Command buttons

The command buttons enable you to perform the following tasks:

Submit

Submits your selections and closes the Select Source Snapshot Copy dialog box.

Cancel

Discards your selections and closes the Select Source Snapshot Copy dialog box.

Related tasks

Resynchronizing protection relationships from the Volume details page on page 301
Resynchronizing protection relationships on page 302

Related references

Resynchronize dialog box on page 347
Initialize/Update dialog box on page 345

Reverse Resync dialog box

When you have a protection relationship that is broken because the source volume is disabled and the destination is made a read/write volume, reverse resynchronization enables you to reverse the direction of the relationship so that the destination becomes the new source and the source becomes the new destination.

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write, while you repair or replace the source, update the source, and reestablish the relationship. When you perform a reverse resync operation, data on the source that is newer than the data on the common Snapshot copy is deleted.

Before reverse resync

Displays the source and destination of a relationship before a reverse resync operation.

Source Volume

The name and location of the source volume before a reverse resync operation .

Destination Volume

The name and location of the destination volume before a reverse resync operation .

After reverse resync

Displays what the source and destination of a relationship is after a reserve resync operation.

Source Volume

The name and location of the source volume after a reverse resync operation.

Destination Volume

The name and location of the destination volume after a reverse resync operation.

Command buttons

The command buttons enable you to perform the following actions:

Submit

Begins the reverse resynchronization process.

Cancel

Closes the Reverse Resync dialog box without initiating a reverse resync operation.

Related tasks

Reversing protection relationships from the Volume details page on page 303
Reversing protection relationships on page 304

Volume Relationships page

The Volume Protection Relationships page displays information about protection objects on the volume, such as transfer and lag details, relationship state, policy, and so on. When the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2, no information is displayed in some columns.

Related tasks

Creating cascade or fanout relationships to extend protection on page 290
Editing protection relationships on page 291
Aborting an active data protection transfer on page 294
Quiescing a protection relationship on page 295
Breaking a SnapMirror relationship on page 297
Removing a protection relationship on page 298
Resuming scheduled transfers on a quiesced relationship on page 298
Resynchronizing protection relationship on page 302
Initializing or updating protection relationships on page 300
Reversing protection relationships on page 304

Related references

Volume Relationships page details on page 351

Volume Protection Relationships page details

The Volume Protection Relationships page displays information about protection relationships on the storage system.

Command buttons

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

Relationship operations

The relationship command buttons enable you to execute the following operations:

- Extend Protection

Enables you to extend protection from an existing relationship by creating either a fanout from the source volume or a cascade from the destination volume of an existing relationship.

- Edit
  Launches the Edit Relationship dialog box, which enables you to change protection policies, schedules, and maximum transfer rates for an existing protection relationship.
- Abort
  Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer.
- Quiesce
  Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress are completed before the relationship is quiesced.
  Break
  Breaks the relationship between the source and destination volume. After a relationship is broken, the destination is made a read-write volume.
- Remove
  Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed.
- Resume
  Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used if one exists.
- Resynchronize
  Enables you to resynchronize a previously broken relationship.
- Initialize/Update
  Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform an update if the relationship is already initialized.
- Reverse Resync
  Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by the contents on the destination.
- Refresh List
  Enables you to refresh the list of resources displayed in the relationship inventory.

Export

Enables you to export the details of all the monitored protection relationships to a comma-separated values (`.csv`) file.

Volume Protection Relationships overview

The Volume Protection Relationships page displays information about protection relationships on the storage system. When the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2, no information is displayed in some columns.

Relationship Status

Displays the current status of the protection relationship. The status can be one of Error (!), Warning (⚠), or Normal (✓).

Lag Status

Displays the lag status for managed relationships, and for unmanaged relationships that have a schedule associated with that relationship. This column is blank when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2. Lag status can be one of the following:

Error (!)

The lag duration is great than or equal to the lag error threshold.

Warning (⚠)

The lag duration is great than or equal to the lag warning threshold.

Normal (✓)

The lag duration is within normal limits.

Transfer Status

Displays the SnapMirror status of a protection relationship. The transfer status can be one of the following:

Idle

Transfers are enabled and no transfer is in progress.

Transferring

SnapMirror transfers are enabled and a transfer is in progress.

Checking

The destination volume is undergoing a diagnostic check and no transfer is in progress.

Quiescing

A SnapMirror transfer is in progress. Additional transfers are disabled.

Quiesced

SnapMirror transfers are disabled. No transfer is in progress.

Queued

SnapMirror transfers are enabled. No transfers are in progress.

Preparing

SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.

Finalizing

SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

Aborting

SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

Relationship Type

Displays the relationship type used to replicate a volume. Relationship types include SnapMirror and SnapVault.

Source SVM

Displays the name of the source SVM.

If the message `Resource-key not discovered` is displayed, this might indicate that the SVM exists on the cluster but has not yet been added to the Unified Manager inventory, or that the SVM was created after the cluster's last refresh. You must ensure that the SVM exists, or perform a rediscovery on the cluster to refresh the list of resources.

You can move your pointer over the source SVM to view information such as the cluster, volume type, protocols allowed, and spaced used. You can view more details about the SVM by clicking on the SVM name.

Source Volume

Displays the source volume being protected. You can view more details about the source volume by clicking the source volume name.

If the message `Resource-key not discovered` is displayed, this might indicate that the volume exists on the cluster but has not yet been added to the Unified Manager inventory, or that the volume was created after the cluster's last refresh. You must ensure that the volume exists, or perform a rediscovery on the cluster to refresh the list of resources.

Destination SVM

Displays the name of the destination SVM.

You can move your pointer over the destination SVM to view information such as the cluster, volume type, protocols allowed, and space used. You can view more details about the SVM by clicking on the SVM name.

Destination Volume

Displays the name of the destination volume.

You can move the pointer over a volume to view information such as the aggregate containing the volume, qtree quota overcommitted space, status of the last volume move operation, and space allocated in the volume. You can also view the details of related objects such as the SVM to which the volume belongs, the aggregate to which the volume belongs, and all the volumes that belong to this aggregate.

Lag Duration

Displays the amount of time that the data on the mirror lags behind the source. This column is blank when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2.

Last Successful Update

Displays the time of the most recent successful data update. The time is displayed in the time zone of the Unified Manager server. This column is hidden by default.

Last Transfer Duration

Displays the time taken for the last data transfer to complete. This column is blank when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2.

Last Transfer Size

Displays the size, in bytes, of the last data transfer. This column is blank when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2.

Relationship Health

Displays the health of the mirror. The health displays Bad if the last manual or scheduled update fails or is aborted, or if the last scheduled update is delayed. The health displays Good if the mirror is healthy and there are no issues. The health column can also remain empty if the health of a mirror is undeterminable. The column is hidden by default.

Relationship State

Displays the current state of the mirror. The state can be one of Uninitialized, Snapmirrored, or Broken-off. The column is hidden by default.

Source Cluster

Displays the source cluster to which the source volume belongs. You can view more details about the source cluster node by clicking the source cluster node name. This column is hidden by default.

Destination Cluster

Displays the cluster that contains the destination volume. You can view more details about the cluster by clicking the cluster name. This column is hidden by default.

Destination Node

Displays the destination node to which the source volume belongs. You can view more details about the destination node by clicking the node name. This column is hidden by default.

Transfer Priority

Displays the priority of the data transfer for a volume. The priority displays either Low or Normal. This column is hidden by default.

Policy

Displays the protection policy for the volume. You can click the policy name to view details associated with that policy, including the following information:

- Transfer priority

  Specifies the priority at which a transfer runs. The transfer priority is Normal or Low. Normal priority transfers are scheduled before low priority transfers. The default is Normal.

- Ignore access time

  Applies only to SnapVault relationships. This specifies whether incremental transfers ignore files which have only their access time changed. The values are either True or False. The default is False.

- Tries limit

  Specifies the maximum number of times to attempt each manual or scheduled transfer for a SnapMirror relationship. The default is 8.

- Comments

  Provides a text field for comments for specific to the selected policy.

- SnapMirror label

  Specifies the SnapMirror label for the first schedule associated with the Snapshot copy policy. The SnapMirror label is used by the SnapVault subsystem when you back up Snapshot copies to a SnapVault destination.

- Retention settings

  Specifies how long backups are kept, based on the time or the number of backups.

- Actual Snapshot copies

  Specifies the number of Snapshot copies on this volume that match the specified label.

- Preserve Snapshot copies

  Specifies the number of SnapVault Snapshot copies that are not deleted automatically even if the maximum limit for the policy is reached. The values are either True or False. The default is False.

- Retention warning threshold

  Specifies the Snapshot copy limit at which a warning is sent to indicate that the maximum retention limit is nearly reached.

Schedule

Displays the name of the protection schedule assigned to the relationship. You can click the schedule name to view details about the schedule.

Version Flexible Replication

Specifies whether Volume Flexible Replication is enabled. The Volume Flexible Replication setting displays either Yes, Yes with backup option, or None. Yes indicates that SnapMirror replication is possible even if source and destination volumes are running different versions of Data ONTAP 8.3 or later. Yes with backup option indicates the implementation of SnapMirror protection with the ability to retain multiple versions of backup copies on the destination. None indicates that Version Flexible Replication is not enabled. This column is hidden by default.

Executing protection workflows using OnCommand Workflow Automation

You can integrate OnCommand Workflow Automation with Unified Manager to execute workflows for your storage classes and monitor SVMs with Infinite Volume that do not have storage classes.

Configuring a connection between OnCommand Workflow Automation and Unified Manager

You can configure a secure connection between Workflow Automation (WFA) and Unified Manager. Connecting to Workflow Automation enables you to use protection features such as SnapMirror and SnapVault configuration workflows, as well as commands for managing SnapMirror relationships.

Before you begin

- You must have the name of the database user that you created in Unified Manager to support WFA and Unified Manager connections.
  This database user must have been assigned the Integration Schema user role.
- You must be assigned either the Administrator role or the Architect role in Workflow Automation.
- You must have the host address, port number 443, user name, and password for the Workflow Automation setup.
- You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Add-ons > Workflow Automation**.
4. In the **Unified Manager Database User** area of the **Set Up OnCommand Workflow Automation** dialog box, select the name, and enter the password for the database user that you created to support Unified Manager and Workflow Automation connections.
5. In the **Workflow Automation Credentials** area of the **Set Up OnCommand Workflow Automation** dialog box, type the host name or IP address (IPv4 or IPv6), and the user name and password for the Workflow Automation setup.

   You must use theUnified Manager server port (port 443).
6. Click **Save and Close**.
7. If you use a self-signed certificate, click **Yes** to authorize the security certificate.

   The Workflow Automation Options Changed dialog box displays.
8. Click **Yes** to reload the web UI, and add the Workflow Automation features.

Related references

Set Up OnCommand Workflow Automation dialog box on page 358

Removing OnCommand Workflow Automation setup from Unified Manager

You can remove the OnCommand Workflow Automation setup from Unified Manager when you no longer want to use Workflow Automation.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click ▦ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Add-ons** > **Workflow Automation**.
4. In the **Set Up OnCommand Workflow Automation** dialog box, click **Remove Setup**.

Related references

Set Up OnCommand Workflow Automation dialog box on page 358

What happens when OnCommand Workflow Automation is reinstalled or upgraded

Before reinstalling or upgrading OnCommand Workflow Automation, you must first remove the connection between OnCommand Workflow Automation and Unified Manager, and ensure that all OnCommand Workflow Automation currently running or scheduled jobs are stopped.

You must also manually delete Unified Manager from OnCommand Workflow Automation.

After you reinstall or upgrade OnCommand Workflow Automation, you must set up the connection with Unified Manager again.

Description of OnCommand Workflow Automation setup windows and dialog boxes

You can set up OnCommand Workflow Automation in Unified Manager by using the Set Up OnCommand Workflow Automation dialog box.

Set Up OnCommand Workflow Automation dialog box

The Set Up OnCommand Workflow Automation dialog box enables you to configure the settings to integrate OnCommand Workflow Automation with Unified Manager. You can also add, modify, or delete the settings.

You must have the OnCommand Administrator or Storage Administrator role.

Unified Manager Database User

This area enables you to enter the credentials of a database user that is required for pairing Unified Manager with OnCommand Workflow Automation:

Name

Enables you to specify the user name of a database user that can be used to access data in the Unified Manager database. By default, no database user is selected. You can select a database user from the drop-down list.

Password

Enables you to specify a password for the specified user name.

OnCommand Workflow Automation Credentials

This area enables you to enter the credentials of an OnCommand Workflow Automation account that is required for pairing with Unified Manager:

Host

Specifies the address of the OnCommand Workflow Automation host server, which is used to pair with Unified Manager.

Port

Specifies the port number of the OnCommand Workflow Automation host server. If the port number is not specified, the default value of the port is 443.

Username

Enables you to specify a user name that can be used to log in to OnCommand Workflow Automation.

Password

Enables you to specify a password for the specified user name.

Command buttons

The command buttons enable you to remove, save, or cancel the setup options:

Remove Setup

Removes the OnCommand Workflow Automation setup from Unified Manager.

Save

Saves the configuration settings for the selected option.

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Configuring a connection between OnCommand Workflow Automation and Unified Manager on page 357

Removing OnCommand Workflow Automation setup from Unified Manager on page 358

Managing and monitoring Infinite Volumes

You can monitor the capacity and availability of your SVMs with Infinite Volume. You can manage the content placement in your SVM with Infinite Volume by creating rules and data policy.

Viewing the details of SVMs with Infinite Volume

You can use the Storage Virtual Machines page to view detailed information about SVMs with Infinite Volume that are monitored by Unified Manager. You can view details such as the capacity, configuration, and data policy and rules associated with the Infinite Volume.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume must be added to the Unified Manager database.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, use the column filter in **Allowed Volume Type** to list the Infinite Volumes that are monitored.
3. View the complete details of the SVM with Infinite Volume by clicking the SVM name.

Related tasks

Adding a user on page 379
Adding an alert on page 102
Editing the Infinite Volume threshold settings on page 361
Editing the threshold settings of storage classes on page 361
Viewing rules on page 366

Viewing the constituents of an Infinite Volume

You can use the Volumes page to view the list of constituents in your Infinite Volume. You can view details such as the constituent state, the SVM with Infinite Volume that contains the constituent, junction path of the constituent, aggregate that contains the constituent, as well as the available, used, and total data capacity of the constituent.

Before you begin

The following requirements must be met:

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume must be added to the Unified Manager database.

Steps

1. Click **Storage > Storage Virtual Machines**.

2. Click the name of the appropriate SVM with Infinite Volume.
3. In the **Storage Virtual Machine details** page, click **Volumes** in the **Related Devices** pane.

   The list of constituents is displayed in the Volumes page.

Related concepts

What a namespace constituent is on page 363
What data constituents are on page 363
What a namespace mirror constituent is on page 364

Related tasks

Adding a user on page 379

Editing the Infinite Volume threshold settings

When you need to address any issues in your Infinite Volume's storage space, you can edit the threshold settings of the Infinite Volume's capacity based on your organization's requirements. When a threshold is crossed, events are generated, and you receive notifications if you have configured alerts for such events.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the required SVM with Infinite Volume.
3. In the **Storage Virtual Machine details** page, click **Actions > Edit Thresholds**.
4. In the **Edit SVM with Infinite Volume Thresholds** dialog box, modify the thresholds as required.
5. Click **Save and Close**.

Related tasks

Adding a user on page 379

Editing the threshold settings of storage classes

When you need to address any issues related to storage space in your storage classes, you can edit the threshold settings of the storage class capacity based on your organization's requirements. When the threshold is crossed, events are generated, and you receive notifications if you have configured alerts for such events.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Storage > Storage Virtual Machines**.

2. In the Storage Virtual Machines page, select the required SVM with Infinite Volume.
3. In the **Storage Virtual Machine details** page, click **Actions** > **Edit Thresholds**.
4. In the **Edit Storage Class Thresholds** dialog box, modify the thresholds as required.
5. Click **Save and Close**.

Related references

Unified Manager user roles and capabilities on page 385

Understanding Infinite Volumes

An Infinite Volume is a logical storage unit that you can use to provide a large, scalable data container with a single namespace and a single mount point. Understanding some of the basic concepts of Infinite Volumes helps you to monitor and manage your SVMs with Infinite Volume.

What an Infinite Volume is

An Infinite Volume is a single, scalable volume that can store up to 2 billion files and tens of petabytes of data.

With an Infinite Volume, you can manage multiple petabytes of data in one large logical entity and clients can retrieve multiple petabytes of data from a single junction path for the entire volume.

An Infinite Volume uses storage from multiple aggregates on multiple nodes. You can start with a small Infinite Volume and expand it nondisruptively by adding more disks to its aggregates or by providing it with more aggregates to use.

Related concepts

Maximum number of files an Infinite Volume can store on page 362

Maximum number of files an Infinite Volume can store

In most cases, an Infinite Volume can hold up to 2 billion files. If an Infinite Volume is relatively small, its maximum number of files might be less than 2 billion.

The maximum number of files that an Infinite Volume can hold is determined by the size of its namespace constituent. If the namespace constituent is 10 TB, the Infinite Volume can hold 2 billion files. If the namespace constituent is less than 10 TB, the Infinite Volume can hold proportionally fewer files.

The size of the namespace constituent is roughly proportional to the size of the Infinite Volume, depending on several factors, such as the namespace constituent's 10 TB maximum size, the available space in the aggregate that holds the namespace constituent, and the SnapDiff setting.

For a two-node Infinite Volume or a multi-node Infinite Volume without SnapDiff enabled, setting the Infinite Volume to a size of 80 TB or greater typically creates a namespace constituent of 10 TB.

The file count not only includes regular files, but also other file system structures, such as directories and symbolic links.

Related concepts

What an Infinite Volume is on page 362

What a storage class is

A storage class is a definition of aggregate characteristics and volume settings. You can define different storage classes and associate one or more storage classes with an Infinite Volume. You must use OnCommand Workflow Automation to define workflows for your storage class requirements and to assign storage classes to Infinite Volumes.

You can define the following characteristics for a storage class:

- Aggregate characteristics, such as the type of disks to use
- Volume settings, such as compression, deduplication, and volume guarantee

For example, you can define a storage class that uses only aggregates with SAS disks and the following volume settings: thin provisioning with compression and deduplication enabled.

The following diagram illustrates an Infinite Volume that spans multiple nodes and uses the following storage classes: gold, silver, and bronze. Each storage class can span two or more nodes within an Infinite Volume. The diagram also illustrates the placement of data constituents in each storage class.

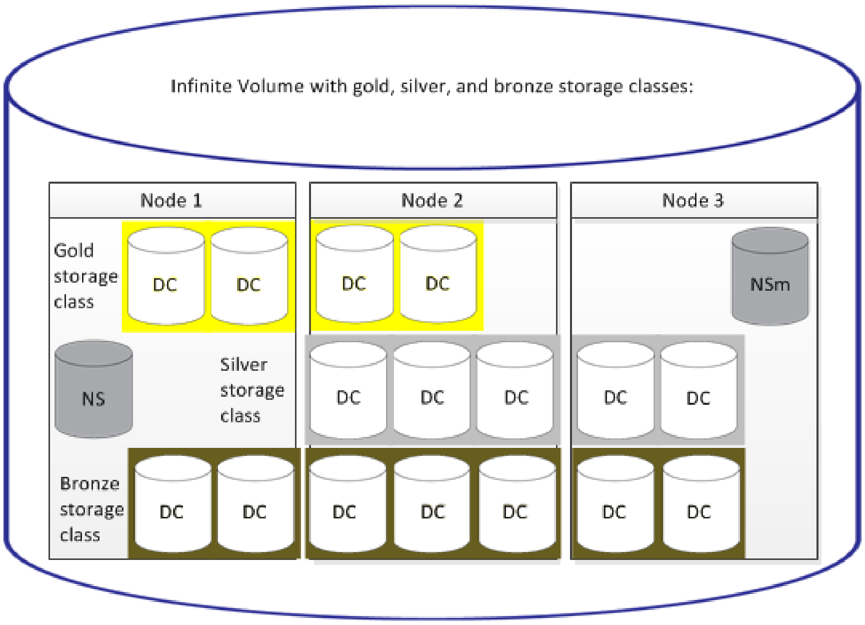

What a namespace constituent is

Each Infinite Volume has a single namespace constituent that maps directory information and file names to the file's physical data location within the Infinite Volume.

Clients are not aware of the namespace constituent and do not interact directly with it. The namespace constituent is an internal component of the Infinite Volume.

What data constituents are

In an Infinite Volume, data is stored in multiple separate data constituents. Data constituents store only the data from a file, not the file's name.

Clients are not aware of data constituents. When a client requests a file from an Infinite Volume, the node retrieves the file's data from a data constituent and returns the file to the client.

Each Infinite Volume typically has dozens of data constituents. For example, a 6 PB Infinite Volume that contains 1 billion files might have 60 data constituents located on aggregates from 6 nodes.

What a namespace mirror constituent is

A namespace mirror constituent is an intracluster data protection mirror copy of the namespace constituent in an Infinite Volume. The namespace mirror constituent performs two roles: It provides data protection of the namespace constituent, and it supports SnapDiff for incremental tape backup of Infinite Volumes.

Creating rules

You can add new rules to your data policy to determine the placement of data that is written to the Infinite Volume. You can create rules either by using rule templates that are defined in Unified Manager or create custom rules.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

Choices

- Creating rules using templates on page 364
- Creating custom rules on page 365

Related tasks

Adding a user on page 379

Creating rules using templates

You can add new rules by using rule templates defined by Unified Manager to determine the placement of data that is written to the Storage Virtual Machine (SVM) with Infinite Volume. You can create rules based on file types, directory paths, or owners.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click the **Data Policy** tab.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Click **Create**.

5. In the **Create Rule** dialog box, choose an appropriate rule template from the drop-down list.

   The template is based on three categories: file type, owner, or directory path.

6. Based on the template selected, add necessary conditions in the **Matching Criteria** area.

7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.

8. Click **Create**.

   The new rule you created is displayed in the Data Policy tab.

9. Optional: Preview any other changes made to the data policy.

10. Click **Activate** to activate the changes in the rule properties in the SVM.

Related concepts

What a rule template is on page 373
What conditions and condition sets are on page 373

Related tasks

Adding a user on page 379

Creating custom rules

Based on your data center requirements, you can create custom rules and add them to a data policy to determine the placement of data that is written to the SVM with Infinite Volume. You can create custom rules from the Create Rule dialog box without using any existing template.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.

2. In the Storage Virtual Machines page, select the appropriate SVM.

3. Click **Data Policy**.

4. Click **Create**.

5. In the **Create Rule** dialog box, select **Custom rule** from the **Template** list.

6. In the **Matching Criteria** area, add conditions as required.

   Conditions enable you to create a rule based on file types, directory paths, or owners. A combination of these conditions are the condition sets. For example, you can have a rule: “Place all .mp3 owned by John in bronze storage class.”

7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.

8. Click **Create**.

   The newly created rule is displayed in the Data Policy tab.

9. Optional: Preview any other changes made to the data policy.

10. Click **Activate** to activate the changes in the rule properties in the SVM.

Related concepts

What conditions and condition sets are on page 373

Related tasks

Adding a user on page 379

Viewing rules

You can view the list of rules you created from the Data Policy tab before modifying the data policy for your SVM with Infinite Volume.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.

Result

The list of rules in the data policy for the selected SVM is displayed. You can use Filter by Storage Class to view rules about a specific storage class.

Related concepts

What rules and data policies are on page 371

Related tasks

Adding a user on page 379
Creating rules on page 364

Editing template-based rules

You can edit a rule that was created using the rule templates from the Edit Rule dialog box. You can add, modify, or delete rule properties such as directory paths, file types, and owners. You can also modify the rule name and the storage class associated with the rule.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate Storage Virtual Machine (SVM).
3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Select the rule for which you want to include new conditions or condition sets.
5. Click **Edit**.
6. In the **Edit Rule** dialog box, edit the rule as required:

If you want to...	Do this...
Add a new rule property	Click **Add**.
Delete a rule property	Click **Delete** by selecting the appropriate rule property.
Modify a rule property	Double-click the appropriate rule property, and then modify as required.

7. Click **Update**.
8. Verify that your modifications are applied to the rule by expanding the rule in the **Data Policy** tab.
9. Optional: Preview any other changes made to the data policy.
10. Click **Activate** to activate the changes to the rule properties in the SVM.

Related concepts

What a rule template is on page 373

Related tasks

Adding a user on page 379

Editing custom rules

You can edit a rule to include new conditions or condition sets in the rule. For example, if you want to include new directory paths along with the owner names, you can do so from the Edit Rule dialog box.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab is visible only for a Storage Virtual Machine (SVM) with Infinite Volume.

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Select the rule for which you want to include new conditions or condition sets.
5. Click **Edit**.
6. In the **Edit Rule** dialog box, add new conditions or condition sets:

If you want to add...	Click...
A new condition	The + icon.
A new condition set	**Add Condition Set**.

7. Click **Update**.
8. Verify that your modifications are applied to the rule by expanding the rule in the **Data Policy** tab.
9. Optional: Preview any other changes made to the data policy.
10. Click **Activate** to activate the changes in the rule properties in the SVM.

Related concepts

What conditions and condition sets are on page 373

Related tasks

Adding a user on page 379

Deleting rules

You can delete a rule from a data policy when it is no longer required. For example, you might want to delete a rule on a particular directory that is no longer valid.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.

3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

4. Select the rule that you want to delete, and then click **Delete**.

   Note: You cannot delete the default rule.

5. Optional: Preview any other changes made to the data policy.

6. Click **Activate** to activate the changes in the rule properties in the SVM.

Related tasks

Adding a user on page 379

Previewing changes to your data policy

You should preview any changes that you have made to your rules in a data policy before you submit the changes in the data policy to the Storage Virtual Machine (SVM) with Infinite Volume for activation.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage** > **Storage Virtual Machines**.

2. In the Storage Virtual Machines page, select the appropriate SVM.

3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

4. Modify the data policy as required.

   Data policy modifications can include creating new rules, editing existing rules, deleting existing rules, or reordering the rules.

5. Click **Activate**.

6. In the **Summary of Changes to Data Policy Configuration** window, preview the changes to your data policy, and then click **Activate** to activate the changes in the data policy in the SVM with Infinite Volume.

Related tasks

Adding a user on page 379

Exporting a data policy configuration

You can export a data policy configuration from Unified Manager to a file. For example, after you have taken the required backup, and in the event of a disaster, you can export the data policy configuration from the primary.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab, which is used while performing this task, is displayed only for SVMs with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Click **Export**.
5. In the browser-specific dialog box, specify the location to which the data policy configuration has to be exported.

Result

The data policy configuration is exported as a JSON file in the specified location.

Related tasks

Importing a data policy configuration on page 370
Adding a user on page 379

Importing a data policy configuration

You can import a data policy configuration from a file, modify the data policy, and then activate the changes in the SVM with Infinite Volume. For example, in the event of a disaster, you can import an already defined data policy to the secondary and modify the policy as required.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you import a data policy configuration, your existing rules are overwritten.

The Data Policy tab is displayed only for SVMs with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.

2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Click **Import**.
5. In the **Import Data Policy** dialog box, specify the data policy that you want to import by providing the absolute path of the data policy file.
6. Click **Import**.
7. Click **Activate** to activate the imported rules in the SVM.

Related tasks

Exporting a data policy configuration on page 370
Adding a user on page 379

Understanding rules and data policy

Understanding the concepts about rules and data policy help you to manage your Infinite Volumes efficiently.

What rules and data policies are

A *rule* determines the placement of files (data) in a Storage Virtual Machine (SVM) with Infinite Volume. A collection of such rules is known as a *data policy*.

Rule

Rules mainly consist of a set of predefined conditions and information that determine where to place files in the Infinite Volume. When a file is placed in the Infinite Volume, the attributes of that file are matched with the list of rules. If attributes match the rules, then that rule's placement information determines the storage class where the file is placed. A default rule in the data policy is used to determine the placement of files if the attributes do not match any of the rules in the rule list.

For example, if you have a rule, “Place all files of type .mp3 in the bronze storage class.”, all .mp3 files that are written to the Infinite Volume would be placed in the bronze storage class.

Data policy

A data policy is a list of rules. Each SVM with Infinite Volume has its own data policy. Each file that is added to the Infinite Volume is compared to its data policy's rules to determine where to place that file. The data policy enables you to filter incoming files based on the file attributes and place these files in the appropriate storage classes.

Related concepts

What the default rule is on page 372
How a data policy filters data written to an Infinite Volume on page 372

Related tasks

Creating rules on page 364

What the default rule is

The default rule is the rule present in the data policy of a Storage Virtual Machine (SVM) with Infinite Volume. It is used to determine the placement of data written to the Infinite Volume when none of the conditions in the existing rules match with the data being written.

The default rule is always the last rule in a data policy and cannot be reordered. For example, consider a data policy with three rules. Rule-1 places all .pdf files in the *high_performance* storage class. Rule-2 places all files owned by the administrator and file names that end with *.xls in the *archival_constituent* storage class. The third rule is the default rule with the *low_performance* storage class.

When a set of *.jpg files that are not owned by the administrator is written to the Infinite Volume, the default rule is used to place these .jpg files in the *low_performance* storage class. Rule-1 and Rule-2 are not used because the data that is written does not match these rules.

Related concepts

What rules and data policies are on page 371

How a data policy filters data written to an Infinite Volume

A data policy automatically filters data written to the Infinite Volume into different storage classes. All files are written to the single file system in the Infinite Volume's namespace, and rules in the data policy determine which storage class stores the data for the files.

A default data policy is automatically created for a Storage Virtual Machine (SVM) with Infinite Volume when you create the Infinite Volume. The data policy is active and contains a default rule. The default rule stores incoming data for files as follows for Infinite Volumes with and without storage classes:

For an Infinite Volume...	The default data policy does this...
Without storage classes	Places all incoming data for files in the Infinite Volume
With one storage class	Places all incoming data for files into the storage class
With one or more storage classes	Places all incoming data for files into the first storage class that is created

Important: For an Infinite Volume with two or more storage classes, you should modify the data policy as soon as possible to create rules that filter data for different types of files into the different storage classes. You should modify the data policy by using OnCommand Unified Manager.

The data policy does not affect the location of the files in the file system in the Infinite Volume's namespace, and storage classes are transparent to client applications. The file system in the namespace contains the file names. The data policy affects only which storage class is used to store the data for the files. Data policies are useful when you assign two or more storage classes to an Infinite Volume.

You can modify the data policy to create additional rules, but you cannot delete the data policy or its default rule.

The following diagram illustrates how a data policy filters data for an Infinite Volume. The file name is stored in the namespace constituent, and rules in the data policy specify that data for this particular file is stored in the silver storage class.

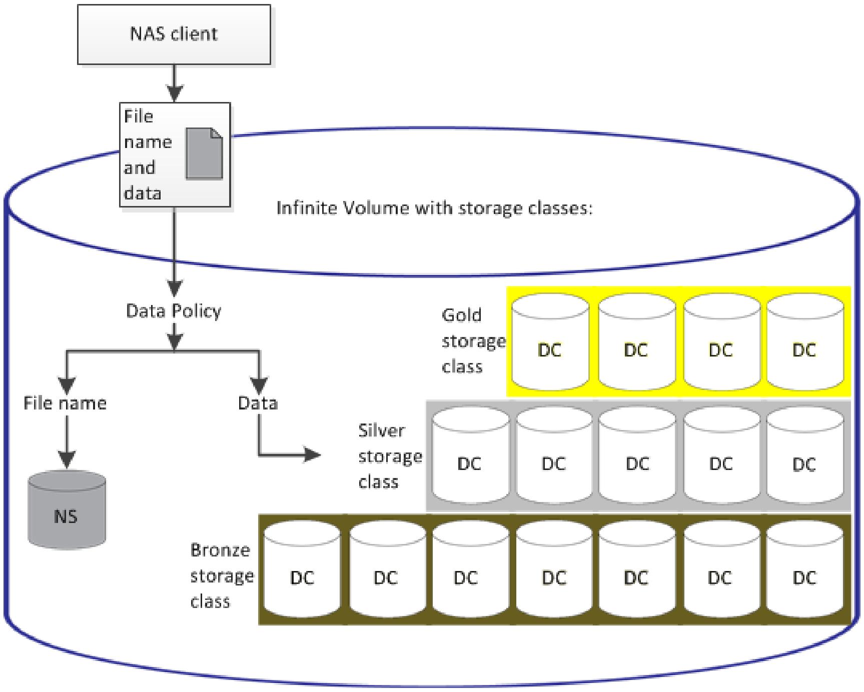

What a rule template is

A rule template is a predefined template that can be used to create rules in a data policy. A rule template enables you to create a rule based on three categories: owner, file type, and directory path.

Example of a rule template for file types

The rule template “Place all files with the specified extensions in a suitable storage class” places all the .mp3 files that are written to the Infinite Volume in a storage class that you specify.

Related concepts

What rules and data policies are on page 371

What conditions and condition sets are

Conditions are a set of matching criteria based on rule properties—such as the file name, directory path, and owner—that define a rule. A collection of such conditions is known as a *condition set*. You can use conditions and condition sets only for custom rules to determine where to place content that is written into your Infinite Volume.

Conditions

For a custom rule, you can specify conditions based on rule properties such as the file name, directory path, or owner, or a combination of all the rule properties. The logic is similar to a Boolean AND operation. For example, by using conditions, you can create a custom rule to place files with `.mp3` extensions and files owned by John in the directory path starting with `/NS/`.

Condition sets

The logic used for condition sets is similar to a Boolean OR operation. For example, by using conditions and condition sets, you can create a complex custom rule that matches either of the following conditions:

- condition-1
  All files owned by Mary and are placed in `/NS/Eng/`
- condition-2
  All files that have names ending with `.pdf` and owned by Mary

Description of Infinite Volume windows and dialog boxes

You can monitor SVMs with Infinite Volume from the respective Storage Virtual Machine details page. You can manage rules and data policies from the Create Rule dialog box. You can also modify the storage class thresholds from the Edit Storage Class Thresholds dialog box.

Create Rule dialog box

You can use the Create Rule dialog box to create new rules for your data policy. For example, when you want to specify the placement of content of a certain file type, you can use the Create Rule dialog box to create the rule for your data policy.

- *Rule Name* on page 374
- *Templates area* on page 374
- *Matching Criteria* on page 374
- *Content Placement* on page 374
- *Command buttons* on page 374

Rule Name

Specifies the name of the new rule.

Templates area

Displays the list of rule templates. You can select an appropriate rule template from the list to create a rule for the data policy.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected. For example, if you select “Place all files with the specified owner names in a suitable storage class”, **List of Owner that...** is displayed in Matching Criteria.

Add

Enables you to add a new rule property based on the rule template selected. For example, if you selected the rule template, “Place all files with the specified owner names in a suitable storage class”, the **Add** button enables you to add the owner's name.

Delete

Enables you to delete a selected rule property.

Content Placement

Enables you to select an appropriate storage class for your rule from the list.

Command buttons

Create

Creates a new rule for the data policy and closes the Create Rule dialog box.

Cancel

Cancels the recent changes made to the rule and closes the Create Rule dialog box.

Related tasks

Creating rules on page 364

Edit Rule dialog box

You can use the Edit Rule dialog box to edit the properties of a rule, such as the file types, directory paths, or owners. You can also select an appropriate storage class for the rule. For example, when a certain file path is no longer valid, you can delete the file path from the corresponding rule.

- *Rule Name* on page 375
- *Matching Criteria* on page 375
- *Content Placement area* on page 375
- *Command buttons* on page 375

Rule Name

Displays the name of the rule.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected.

Add

Enables you to add a new rule property, a new file type, a file path, or a new owner. For example, if you had specified the rule template, “Place all files with the specified owner names in a suitable storage class”, the Add button enables you to add the owner's name.

Delete

Enables you to delete a selected rule property.

Content Placement area

Displays the list of storage classes. You can select an appropriate storage class for the selected rule.

Command buttons

Update

Updates the changes made to the rule and closes the Edit Rule dialog box.

Cancel

Cancels the recent changes made to the rule and closes the Edit Rule dialog box.

Related tasks

Editing template-based rules on page 367
Editing custom rules on page 367

Edit Rule dialog box (Advanced edit)

You can use the Edit Rule dialog box to edit the properties of a rule that is not created by using a template. The rule properties you can edit include the file types, directory paths, matching criteria, or owners. You can select an appropriate storage class for the rule. For example, you can edit the conditions specified in the matching criteria of a rule.

- *Rule Name* on page 376
- *Matching Criteria* on page 376
- *Content Placement area* on page 376
- *Command buttons* on page 376

Rule Name

Displays the name of the rule.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected. You can expand the rules and modify the rule properties, as required.

Content Placement area

Displays the list of storage classes. You can select an appropriate storage class for the selected rule.

Command buttons

Update

Updates the changes made to the rule and closes the Edit Rule dialog box.

Cancel

Cancels the recent changes made to the rule and closes the Edit Rule dialog box.

Edit SVM with Infinite Volume Thresholds dialog box

You can use the Edit SVM with Infinite Volume Thresholds dialog box to modify the default threshold values of each SVM with Infinite Volume, based on your organization's requirements. The default threshold values indicate the level of activity that must be reached on the SVM before an event is triggered.

Capacity

The Capacity area enables you to set capacity threshold conditions for the selected SVM with Infinite Volume:

Space Nearly Full

Specifies the percentage at which the SVM with Infinite Volume is considered to be nearly full. It also displays the corresponding space (in GB, MB, or TB) in the Infinite Volume. For example, if you have an Infinite Volume of size 10 GB, and the Space Nearly Full threshold is 80%, then the following information is displayed: (8 GB of 10 GB).

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which the SVM with Infinite Volume is considered full. It also displays the corresponding space (in GB, MB, or TB) in the Infinite Volume. For example, if you have an Infinite Volume of size 10 GB, and the Space Full threshold is 90%, then the following information is displayed: (9 GB of 10 GB).

You can also use the slider to set the threshold value.

Snapshot Usage Limit

Specifies the limit, in percentage, of space reserved for Snapshot copies in the Infinite Volume.

Command buttons

The command buttons enable you to perform the following tasks:

Restore to Global Defaults

Enables you to restore the threshold settings to the current values that are set at the global level.

Save

Saves all the threshold settings.

Save and Close

Saves all the threshold settings and then closes the Edit SVM with Infinite Volume Thresholds dialog box.

Cancel

Ignores any changes to the threshold settings and closes the Edit SVM with Infinite Volume Thresholds dialog box.

Related tasks

Editing the Infinite Volume threshold settings on page 361

Edit Storage Class Thresholds dialog box

You can use the Edit Storage Class Thresholds dialog box to modify the default threshold values of various storage classes in each SVM with Infinite Volume based on your organization's requirements. The default threshold values indicate the level of activity that must be reached on a storage class before an event is triggered.

You must have the OnCommand Administrator or Storage Administrator role.

Capacity

The Capacity area enables you to set capacity threshold conditions for the selected storage class.

Space Nearly Full

Specifies the percentage at which a storage class in the SVM with Infinite Volume is considered to be nearly full. It also displays the corresponding space (in GB, MB, or TB) in the storage class. For example, if you have a storage class of size 10 GB and the Space Nearly Full threshold is 80%, then the following information is displayed: (8 GB of 10 GB).

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which the storage class in the SVM with Infinite Volume is considered full. It also displays the corresponding space (in GB, MB, or TB) in the storage class. For example, if you have a storage class of size 10 GB and the Space Full threshold is 90%, then the following information is displayed: (9 GB of 10GB).

You can also use the slider to set the threshold value.

Snapshot Usage Limit

Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.

Command buttons

The command buttons enable you to perform tasks for a selected volume.

Restore to Global Defaults

Enables you to restore the threshold settings to the current values that are set at the global level.

Save

Saves all the threshold settings.

Save and Close

Saves all the threshold settings and then closes the Edit Storage Class Thresholds dialog box.

Cancel

Cancels changes (if any) to the threshold settings and closes the Edit Storage Class Thresholds dialog box.

Related tasks

Editing the threshold settings of storage classes on page 361

Managing user access

You can create roles and assign capabilities to control user access to selected cluster objects. You can identify users who have the required capabilities to access selected objects within a cluster. Only these users are provided access to manage the cluster objects.

Adding a user

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

Before you begin

- To add a remote user or group, you must have enabled remote authentication and configured your authentication server.
- You must have the OnCommand Administrator role.

About this task

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

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Users**.
3. On the **Manage Users** page, click **Add**.
4. In the **Add User** dialog box, select the type of user that you want to add, and enter the required information.

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

Related tasks

Enabling remote authentication on page 395
Setting up authentication services on page 397
Adding authentication servers on page 398

Related references

Unified Manager user roles and capabilities on page 385
Definitions of user types on page 383
Definitions of user roles in Unified Manager on page 384

Editing the user settings

You can edit user settings—such as the email address and role—that are specified each user. For example, you might want to change the role of a user who is a storage operator, and assign storage administrator privileges to the user.

Before you begin

You must have the OnCommand Administrator role.

About this task

When you modify the role that is assigned to a user, the changes are applied when either of the following actions occur:

- The user logs out and logs back in to Unified Manager.
- Session timeout of 24 hours is reached.

Steps

1. Click > **Health**.
2. Click **Administration** > **Manage Users**.
3. In the **Manage Users** page, select the user for which you want to edit settings, and click **Edit**.
4. In the **Edit User** dialog box, edit the appropriate settings that are specified for the user.
5. Click **Save**.

Related tasks

Adding a user on page 379

Testing a remote user or remote group

You can validate that a remote user or remote group can access the Unified Manager server by using the authentication settings that are specified for your authentication servers.

Before you begin

- You must have enabled remote authentication, and configured your authentication settings so that the Unified Manager server can validate the remote user or remote group.
- You must have the OnCommand Administrator role.

Steps

1. Click > **Health**.
2. Click **Administration** > **Manage Users**.
3. In the **Manage Users** page, select a remote user or remote group that you want to validate, and then click **Test**.

Related tasks

Adding a user on page 379
Enabling remote authentication on page 395
Setting up authentication services on page 397
Testing the configuration of authentication servers on page 399

Viewing users

You can use the Manage Users page to view the list of users who manage storage objects and data using Unified Manager. You can view details about the users, such as the user name, type of user, email address, and the role that is assigned to the users.

Before you begin

You must have the OnCommand Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Manage Users**.

   The list of users is displayed in the Manage Users page.

Related tasks

Adding a user on page 379

Deleting users or groups

You can delete one or more users from the management server database to prevent specific users from accessing Unified Manager. You can also delete groups so that all the users in the group can no longer access the management server.

Before you begin

- When you are deleting remote groups, you must have reassigned the events that are assigned to the users of the remote groups.
  If you are deleting local users or remote users, the events that are assigned to these users are automatically unassigned.
- You must have the OnCommand Administrator role.

Steps

1. Click **Health** > **Administration** > **Manage Users**.
2. In the **Manage Users** page, select the users or groups that you want to delete, and then click **Delete**.
3. Click **Yes** to confirm the deletion.

Related tasks

Adding a user on page 379
Assigning events on page 78

Changing the local user password

You can change your login password to prevent potential security risks. If you have configured Unified Manager in a VCS environment, then you must change the password for both cluster nodes. Both cluster nodes must have same password.

Before you begin

You must be logged in as a local user.

About this task

The passwords for the maintenance user and for the remote user cannot be changed from the web UI. You must use the Unified Manager maintenance console to change the maintenance user password. To change the remote user password, you must contact your password administrator.

Steps

1. Log in to Unified Manager.
2. Click ***`user_name`*** > **Change Password**.

   The **Change Password** option is not displayed if you are a remote user.
3. In the **Change Password** dialog box, enter the details as required.
4. Click **Save**.

Related tasks

Adding a user on page 379

What the maintenance user does

If Unified Manager is installed as a virtual appliance, the maintenance user is created during initial configuration, the maintenance user can create subsequent users and assign them roles. The maintenance user has OnCommand administrator role in the web UI. If Unified Manager is installed as a virtual appliance, the maintenance user can also access the Unified Manager maintenance console.

If Unified Manager is installed as a virtual appliance, the maintenance user can perform the following functions using the maintenance console:

- Configure network access
- Upgrade to newer versions of Unified Manager
- Shut down virtual appliances (only from VMware console)
- Increase data disk or swap disk size
- Change the time zone
- Generate support bundles to send to technical support

Note: If Unified Manager is installed on Red Hat Enterprise Linux, no maintenance console is supplied and no maintenance user is created. Equivalent operations are instead performed by the Red hat Enterprise Linux root user using Linux command lines.

Related references

Definitions of user types on page 383
Definitions of user roles in Unified Manager on page 384
Unified Manager user roles and capabilities on page 385

What RBAC is

RBAC (role-based access control) provides the ability to control who has access to various features and resources in the OnCommand Unified Manager server.

What role-based access control does

Role-based access control (RBAC) enables administrators to manage groups of users by defining roles. If you need to restrict access for specific functionality to selected administrators, you must set up administrator accounts for them. If you want to restrict the information that administrators can view and the operations they can perform, you must apply roles to the administrator accounts you create.

The management server uses RBAC for user login and role permissions. If you have not changed the management server's default settings for administrative user access, you do not need to log in to view them.

When you initiate an operation that requires specific privileges, the management server prompts you to log in. For example, to create administrator accounts, you must log in with Administrator account access.

Definitions of user types

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

Unified Manager user types are as follows:

Maintenance user

Created during the initial configuration of Unified Manager. The maintenance user then creates additional users and assigns roles. The maintenance user is also the only user with access to the maintenance console. If Unified Manager is installed as a virtual appliance, the maintenance user is also the only user with access to the maintenance console. If Unified Manager is installed on Red Hat Enterprise Linux, the maintenance user is given the user name "umadmin."

Local user

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

Remote group

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

Remote user

Accesses the Unified Manager UI using the credentials stored on the authentication server. A remote user performs functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.

Database user

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

Related concepts

What the maintenance user does on page 382

Definitions of user roles in Unified Manager

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

Unified Manager includes the following predefined user roles:

Operator

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

Storage Administrator

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

OnCommand Administrator

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

Note: If Unified Manager is installed on Red Hat Enterprise Linux, the initial user with the OnCommand Administrator role is automatically named "umadmin." If installed on Red Hat Enterprise Linux, networking administration is not supported in the Unified Manager web UI.

Event Publisher

Transmits events generated by partner applications for display in the Unified Manager web UI. This specialized, limited-access user role enables partner applications to share event information with Unified Manager. At the same time, the limited access of this role prevents unauthorized access to the Unified Manager server or the Unified Manager web UI through event publishing activity.

Integration Schema

The role enables read-only access to Unified Manager database views for integrating Unified Manager with WFA.

Report Schema

The role enables read-only access to report-specific database views directly from the database.

Related concepts

What the maintenance user does on page 382

Unified Manager user roles and capabilities

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

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

Function	**Event Publisher**	**Operator**	**Storage Administrator**	**OnCommand Administrator**	**Integration Schema**	**Report Schema**
View storage system information		•	•	•		
View other data, such as histories and capacity trends		•	•	•		
View, assign, and resolve events		•	•	•		
View storage service objects, such as SVM associations and resource pools		•	•	•		
Manage storage service objects, such as SVM associations and resource pools			•	•		
Define alerts			•	•		
Manage storage management options			•	•		
Manage storage management policies			•	•		
Manage users				•		
Manage administrative options				•		
Manage database access				•		
Publish events	•					

Function	Event Publisher	Operator	Storage Administrator	OnCommand Administrator	Integration Schema	Report Schema
Manage integration with WFA and provide access to the database views					•	
Provide access to the report database views						•
Schedule and save reports		•	•	•		
Import and delete imported reports				•		

Related concepts

What the maintenance user does on page 382

Description of user access windows and dialog boxes

Based on the RBAC settings, you can add users from the Manage Users page and assign appropriate roles to those users to access and monitor your clusters.

Manage Users page

The Manage Users page displays a list of your users and groups, and provides information such as the name, type of user, and email address. You can also use this page to perform tasks such as adding, editing, deleting, and testing users.

- *Command buttons* on page 386
- *List view* on page 387

Command buttons

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

Add

Displays the Add User dialog box, which enables you to add a local user, a remote user, a remote group, or a database user.

You can add remote users or groups only if your authentication server is enabled and configured.

Edit

Displays the Edit User dialog box, which enables you to edit the settings for the selected user.

Delete

Deletes the selected users from the management server database.

Test

Enables you to validate whether a remote user or group is present in the authentication server.

You can perform this task only if your authentication server is enabled and configured.

List view

The List view displays, in tabular format, information about the users that are created. You can use the column filters to customize the data that is displayed.

Name

Displays the name of the user or group.

Type

Displays the type of user: Local User, Remote User, Remote Group, Database User, or Maintenance User.

Email

Displays the email address of the user.

Role

Displays the type of role that is assigned to the user: Operator, Storage Administrator, OnCommand Administrator, Event Publisher, Integration Schema, or Report Schema.

Related tasks

Viewing users on page 381
Adding a user on page 379
Editing the user settings on page 380
Testing a remote user or remote group on page 380
Deleting users or groups on page 381

Add User dialog box

You can create local users or database users, or add remote users or remote groups, and assign roles so that these users can efficiently manage storage objects and data using Unified Manager.

You can add a user by completing the following fields:

Type

Enables you to specify the type of user you want to create.

Name

Enables you to specify a user name that a user can use to log in to Unified Manager.

Password

Enables you to specify a password for the specified user name. This field is displayed only when you are adding a local user or a database user.

Confirm Password

Enables you to reenter your password to ensure the accuracy of what you entered in the Password field. This field is displayed only when you are adding a local user or a database user.

Email

Enables you to specify an email address for the user; the email address specified must be unique to the user name. This field is displayed only when you are adding a remote user or a local user.

Role

Enables you to assign a role to the user and defines the scope of activities that the user can perform. The role can be OnCommand Administrator, Storage Administrator, Operator, Event Publisher, Integration Schema, or Report Schema.

Command buttons

The command buttons enable you to perform the following tasks:

Add

Adds the user and closes the Add User dialog box.

Cancel

Cancels the changes and closes the Add User dialog box.

Related tasks

Adding a user on page 379

Edit User dialog box

The Edit User dialog box enables you to edit only certain settings, depending on the selected user.

Details

The Details area enables you to edit the following information about a selected user:

Type

This field cannot be edited.

Name

This field cannot be edited.

Password

Enables you to edit the password when the selected user is a database user.

Confirm Password

Enables you to edit the confirmed password when the selected user is a database user.

Email

Enables you to edit the email address of the selected user. This field can be edited when the selected user is a local user, LDAP user, or maintenance user.

Role

Enables you to edit the role that is assigned to the user. This field can be edited when the selected user is a local user, remote user, or remote group.

Command buttons

The command buttons enable you to perform the following tasks:

Save

Saves the changes and closes the Edit User dialog box.

Cancel

Cancels the changes and closes the Edit User dialog box.

Related tasks

Editing the user settings on page 380

Configuring backup and restore operations

You can create scheduled backups of Unified Manager and use the restore feature to restore the backup to a local system or a remote system.

What database backup is

A backup is a copy of the Unified Manager database and configuration files that you use in case of a system failure or data loss. In Unified Manager, you can perform a scheduled local or remote backup.

You can create a scheduled backup by adding or editing the backup setting attributes. By default, the scheduled backup is disabled. You can also view backup failure and success events.

Before beginning a backup operation, Unified Manager performs an integrity check to verify that all the required backup files and backup directories exist and are writable.

You can restore a Unified Manager backup only on the same version of Unified Manager. For example, if you created a backup on Unified Manager 6.3, the backup can be restored only on Unified Manager 6.3.

Related tasks

Configuring database backup settings on page 389
Configuring database backup settings on page 389

Configuring database backup settings

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

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- Verify that the “jboss” user has write permissions to the backup directory.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Database Backup**.
3. In the **Backup and Restore** page, click **Actions** > **Database Backup Settings**.
4. Configure the appropriate values for a backup path and retention count.

   The default value for retention count is 10; you can use 0 for creating unlimited backups.
5. Select **Schedule Frequency**.
6. In the **Backup Schedule** section, specify a daily or weekly schedule.

   Daily

   If you select this option, you must enter a time in 24-hour format for creating the backup. For example, if you specify 18:30, then a backup is created daily at 6:30 PM.

Weekly

If you select this option, you must specify the time and day for creating the backup. For example, if you specify the day as Monday and time as 16:30, then a weekly backup is created every Monday at 4:30 PM.

7. Click **Save and Close**.

Related concepts

What database backup is on page 389
What database backup is on page 389

What a database restore is

Database restore is the process of copying backup files from secondary storage to a disk to restore the original data if data loss occurs. You perform the restore operation from the Unified Manager console.

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

The restore operation is performed using restore commands that are executed on the Unified Manager server from the console. The restore feature provides status messages related to restore success or failure.

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

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

Important: Old backup files cannot be used to restore an image after Unified Manager has been upgraded to a newer version of software. These old backup files are not removed automatically based on the retention count value you have set. The retention count setting only applies to the backup files created with the current version of Unified Manager software. Therefore, you must manually delete these old backup files from the backup folder so that they do not waste space on the system.

Related tasks

Restoring a database backup on Red Hat Enterprise Linux on page 390
Restoring a database backup on a virtual machine on page 391
Restoring a database backup on Windows on page 392

Restoring a database backup on Red Hat Enterprise Linux

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

Before you begin

- You must have the root user credentials for the Red Hat Enterprise Linux host on which Unified Manager is installed.

- You must have installed and configured Unified Manager.
- The Unified Manager backup file must already exist on the system on which you will perform the restore operation.
- The backup file must be of `7z` type.

About this task

If Unified Manager has been upgraded from version 6.3 and earlier, the database backup files are located in `/data/ocum-backup`. If Unified Manager has been installed new with version 6.4 or later, or upgraded from version 6.4 or later, the database backup files are located in `/opt/netapp/data/ocum-backup`.

Steps

1. Log in as the root user to the Red Hat Enterprise Linux host on which Unified Manager is installed.
2. If Unified Manager is installed in VCS setup, then stop the Unified Manager ocie, ocieau, and rp services using Veritas Operations Manager.
3. At the command prompt, restore the backup:

   `um backup restore -f /opt/netapp/data/ocum-backup/`*`backup_file_name`*

   Example

   ```
   um backup restore -f /opt/netapp/data/ocum-backup/
   UM_6.4.N151113.1348_backup_unix_11-25-2015-04-45.7z
   ```

   If the folder name contains a space, you must include the absolute path or relative path in double quotation marks; for example: **`"/opt/netapp/data/ocum-backup/UM_6.4.N151118.2300_backup_rhel_11-20-2015-02-51.7z"`**

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

Related concepts

What a database restore is on page 390

Restoring a database backup on a virtual machine

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

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have installed and configured Unified Manager.
- You must have copied a Unified Manager backup file to the system on which you want to perform the restore operation.
- The backup file must be of `.7z` type.

About this task

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

Steps

1. In the vSphere client, locate the Unified Manager virtual machine, and then select the **Console** tab.
2. Click in the console window, and then log in to the maintenance console using your user name and password.
3. In the **Main Menu**, enter the number for the **System Configuration** option.
4. In the **System Configuration Menu**, enter the number for the **Restore Database from bundle** option.
5. When prompted, enter the absolute path of the backup file.

   #### Example

   ```
   Bundle to retore from: opt/netapp/data/ocum-backup/
   UM_6.4.N151112.0947_backup_unix_11-22-2015-11-41.7z
   ```

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

After you finish

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

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

Note: When OnCommand Performance Manager is paired with a Unified Manager server, and you restore the backup to a different system, you must update Performance Manager with the new IP address for the Unified Manager server.

Related concepts

What a database restore is on page 390

Restoring a database backup on Windows

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

Before you begin

- You must have Windows administrator privileges.
- You must have installed and configured Unified Manager.
- The Unified Manager backup file that you want to restore must exist in the system on which you want to perform the restore operation.

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

About this task

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

Steps

1. Log in to the Unified Manager console as an administrator:

   ```
   um cli login -u maint_username
   ```

2. At the command prompt, restore the backup:

   ```
   um backup restore -f \ProgramData\NetApp\OnCommandAppData\ocum\backup
   \backup_file_name
   ```

   Example

   ```
   um backup restore -f \ProgramData\NetApp\OnCommandAppData\ocum\backup
   \UM_6.4.N151118.2300_backup_windows_11-20-2015-02-51.7z
   ```

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

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

Related concepts

What a database restore is on page 390

Description of backup windows and dialog boxes

You can view the list of backups from the backup page in Unified Manager. You can view the backup name, size, and creation time for the backups listed in this page. You can modify the database backup settings from the Database Backup Settings page.

Database Backup and Restore page

The Backup and Restore page displays a list of backups created by Unified Manager and provides information about the backup name, size, creation time, and schedule. You can also restore backups from this page.

You must have the OnCommand Administrator or Storage Administrator role.

Command buttons

Actions

Displays the Database Backup Settings dialog box, which enables you to specify a backup path, retention count, and backup schedule.

List View

The list view displays, in tabular format, information about the backups created by Unified Manager. You can use the column filters to customize the data that is displayed.

Name

Displays the name of the selected backup.

Size

Displays the size of the selected backup.

Creation Time

Displays the creation date and time of the selected backup.

Schedule

Displays the status of the backup and restore operation. Also indicates whether it is a scheduled backup or not.

Database Backup Settings dialog box

You can use the Database Backup Settings dialog box to specify a backup path and retention count and to enable a backup schedule for a selected backup instance.

You can change the following database backup settings:

Path

Specifies the path to the location where you store the backup files. The following table specifies the backup path format, and default locations, for different operating systems:

Host operating system	Backup path format
Virtual application	`/opt/netapp/data/backup`
Red Hat Enterprise Linux	`/opt/netapp/data/backup`
Microsoft Windows	`C:\ProgramData\NetApp\OnCommandAppData\ocum\backup\`

Retention Count

Specifies the maximum number of backups to be retained by Unified Manager. The default value is ten.

Schedule Frequency Enable

This option enables you to specify when to schedule a backup; you can choose daily or weekly.

Daily

Specifies the daily backup schedule with the time.

Weekly

Specifies the weekly backup schedule with the day and time.

Command buttons

Save and Close

Saves the backup file and closes the dialog box. Unified Manager saves the backup file in the following format: `um_um_version_backup_os_timestamp.7z`.

Cancel

Closes the Database Backup Settings dialog box without saving your changes.

Managing authentication

You can enable LDAP authentication on the Unified Manager server and configure it to work with your LDAP servers to authenticate remote users.

Enabling remote authentication

You can enable remote authentication by using either Open LDAP or Active Directory, so that the management server can communicate with your authentication servers. The users of the authentication server can use Unified Manager to manage storage objects and data.

Before you begin

You must have the OnCommand Administrator role.

About this task

If remote authentication is disabled, remote users cannot access Unified Manager.

Remote authentication is supported over LDAP and LDAPS (Secure LDAP). Unified Manager uses 389 as the default port for non-secure communication, and 636 as the default port for secure communication.

Steps

1. Click ▦▾ **> Health**.
2. Click **Administration > Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server > Authentication**.
4. Select **Enable Remote Authentication**.
5. In the **Authentication Service** field, select **Active Directory** or **Open LDAP**.
6. Configure the authentication service.

For Authentication type...	Enter the following information...
Active Directory	• Authentication server administrator name in one of following formats: ◦ *`domainname\username`* ◦ *`username@domainname`* ◦ *`Bind Distinguished Name`* (using the appropriate LDAP notation) • Administrator password • Base distinguished name (using the appropriate LDAP notation)

For Authentication type...	Enter the following information...
Open LDAP	• Bind distinguished name (in the appropriate LDAP notation) • Bind password • Base distinguished name

If the authentication of an Active Directory user takes a long time or times out, the authentication server is probably taking a long time to respond. Disabling support for nested groups in Unified Manager might reduce the authentication time.

If you select the Use Secure Connection option for the authentication server, then Unified Manager communicates with the authentication server using the Secure Sockets Layer (SSL) protocol.

7. Optional: Add authentication servers, and test the authentication.
8. Click **Save and Close**.

Related tasks

Adding a user on page 379

Disabling nested groups from remote authentication

If you have remote authentication enabled, you can disable nested group authentication so that only individual users and not group members can remotely authenticate to Unified Manager. You can disable nested groups when you want to improve Active Directory authentication response time.

Before you begin

You must be logged in as an Active Directory domain user to perform this task. Logging in as an Active Directory administrator is not required.

About this task

Disabling support for nested groups in Unified Manager might reduce the authentication time. If nested group support is disabled, and if a remote group is added to Unified Manager, individual users must be members of the remote group to authenticate to Unified Manager.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > Authentication**.
3. Select **Enable Remote Authentication**.
4. In the **Authentication Service** field, select **Others**.
5. In the **Member** field, change the member information from “member:1.2.840.113556.1.4.1941:” to “member”.
6. Click **Save and Close**.

Setting up authentication services

Authentication services enable the authentication of remote users or remote groups in an authentication server before providing them access to Unified Manager. You can authenticate users by using predefined authentication services (such as Active Directory or OpenLDAP), or by configuring your own authentication mechanism.

Before you begin

- You must have enabled remote authentication.
- You must have the OnCommand Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration > Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server > Authentication**.
4. Select one of the following authentication services:

If you select...	Then do this...
Active Directory	**a.** Enter the administrator name and password. **b.** Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **`cn=ou,dc=domain,dc=com`**.
OpenLDAP	**a.** Enter the bind distinguished name and bind password. **b.** Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **`cn=ou,dc=domain,dc=com`**.
Others	**a.** Enter the bind distinguished name and bind password. **b.** Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **`cn=ou,dc=domain,dc=com`**. **c.** Specify the LDAP protocol version that is supported by the authentication server. **d.** Enter the user name, group membership, user group, and member attributes.

Note: If you want to modify the authentication service, you must delete any existing authentication servers, and then add new authentication servers.

5. Click **Save and Close**.

Related tasks

Adding a user on page 379
Enabling remote authentication on page 395

Adding authentication servers

You can add authentication servers and enable remote authentication on the management server to enable remote users within the authentication server to access Unified Manager.

Before you begin

- The following information must be available:
  - Host name or IP address of the authentication server
  - Port number of the authentication server
- You must have enabled remote authentication and configured your authentication service so that the management server can authenticate remote users or groups in the authentication server.
- You must have the OnCommand Administrator role.

About this task

If the authentication server that you are adding is part of a high-availability (HA) pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.

Steps

1. Click ▦ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server** > **Authentication**.
4. Enable or disable the **Use secure connection authentication** option:

If you want to...	Then do this...
Enable it	**a.** In Enable Remote Authentication area, select the **Use Secure Connection** option. **b.** In the Servers area, click **Add**. **c.** In the Add Authentication Server dialog box, enter the authentication name or IP address (IPv4 or IPv6) of the server. **d.** In the Authorize Host dialog box, click View Certificate. **e.** In the View Certificate dialog box, verify the certificate information, and then click **Close**. **f.** In the Authorize Host dialog box, click **Yes**. **Note:** When you enable the **Use Secure Connection authentication** option, Unified Manager communicates with the authentication server and displays the certificate. Unified Manager uses 636 as default port for secure communication and port number 389 for non-secure communication.

If you want to...	Then do this...
Disable it	**a.** In the Enable Remote Authentication area, clear the **Use Secure Connection** option. **b.** In the Servers area, click **Add**. **c.** In the Add Authentication Server dialog box, specify either the host name or IP address (IPv4 or IPv6) of the server, and the port details. **d.** Click **Add**.

The authentication server that you added is displayed in the Servers area.

5. Perform a test authentication to confirm that you can authenticate users in the authentication server that you added.

Related concepts

Authentication with Active Directory or OpenLDAP on page 401

Related tasks

Adding a user on page 379
Enabling remote authentication on page 395
Setting up authentication services on page 397
Testing the configuration of authentication servers on page 399

Testing the configuration of authentication servers

You can validate the configuration of your authentication servers to ensure that the management server is able to communicate with them. You can validate the configuration by searching for a remote user or remote group from your authentication servers, and authenticating them using the configured settings.

Before you begin

- You must have enabled remote authentication, and configured your authentication service so that the Unified Manager server can authenticate the remote user or remote group.
- You must have added your authentication servers so that the management server can search for the remote user or remote group from these servers and authenticate them.
- You must have the OnCommand Administrator role.

About this task

If the authentication service is set to Active Directory, and if you are validating the authentication of remote users who belong to the primary group of the authentication server, information about the primary group is not displayed in the authentication results.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server** > **Authentication**.

4. In the **Authentication Setup Options** dialog box, click **Test Authentication**.
5. In the **Test User** dialog box, specify the user name and password of the remote user or the user name of the remote group, and then click **Test**.

   If you are authenticating a remote group, you must not enter the password.

Related tasks

Adding a user on page 379
Enabling remote authentication on page 395
Setting up authentication services on page 397
Adding authentication servers on page 398

Editing authentication servers

You can change the port that the Unified Manager server uses to communicate with your authentication server.

Before you begin

You must have the OnCommand Administrator role.

Steps

1. Click ▦ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server** > **Authentication**.
4. In the **Servers** area, select the authentication server that you want to edit, and then click **Edit**.
5. In the **Edit Authentication Server** dialog box, edit the port details.
6. Click **Save**.

Related tasks

Adding a user on page 379

Deleting authentication servers

You can delete an authentication server if you want to prevent the Unified Manager server from communicating with the authentication server. For example, if you want to change an authentication server that the management server is communicating with, you can delete the authentication server and add a new authentication server.

Before you begin

You must have the OnCommand Administrator role.

About this task

When you delete an authentication server, remote users or groups of the authentication server will no longer be able to access Unified Manager.

Steps

1. Click **Health > Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > Authentication**.
3. In the **Servers** area, select one or more authentication servers that you want to delete, and then click **Delete**.
4. Click **Yes** to confirm the delete request.

   If the **Use Secure Connection** option is enabled, then the certificates associated with the authentication server are deleted along with the authentication server.

Related tasks

Adding a user on page 379

Authentication with Active Directory or OpenLDAP

You can enable remote authentication on the management server and configure the management server to communicate with your authentication servers so that users within the authentication servers can access Unified Manager.

You can use one of the following predefined authentication services or specify your own authentication service:

- Microsoft Active Directory

  Note: You cannot use Microsoft Lightweight Directory Services.

- OpenLDAP

You can select the required authentication service and add the appropriate authentication servers to enable the remote users in the authentication server to access Unified Manager. The credentials for remote users or groups are maintained by the authentication server. The management server uses the Lightweight Directory Access Protocol (LDAP) to authenticate remote users within the configured authentication server.

For local users who are created in Unified Manager, the management server maintains its own database of user names and passwords. The management server performs the authentication and does not use Active Directory or OpenLDAP for authentication.

Description of authentication windows and dialog boxes

You can enable LDAP authentication from the Authentication Setup Options dialog box.

Authentication Setup Options dialog box

You can use the Authentication Setup Options dialog box to configure the management server to communicate with your authentication server and authenticate remote users in the authentication server.

You must have the OnCommand Administrator or Storage Administrator role.

Enable Remote Authentication area

The Enable Remote Authentication area allows you to enable or disable remote authentication. You can enable remote authentication to enable the management server to authenticate remote users within the configured authentication servers.

Authentication Service

Enables you to configure the management server to authenticate users in directory service providers, such as Active Directory, OpenLDAP, or specify your own authentication mechanism. You can specify an authentication service only if you have enabled remote authentication.

Active Directory

- Administrator Name
  Specifies the administrator name of the authentication server.
- Password
  Specifies the password to access the authentication server.
- Base Distinguished Name
  Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **`dc=ou,dc=domain,dc=com`**.
- Disable Nested Group Lookup
  Specifies whether to enable or disable the nested group lookup option. By default this option is disabled. If you use Active Directory, you can speed up authentication by disabling support for nested groups.
- Use Secure Connection
  Specifies the authentication service used for communicating with authentication servers.

OpenLDAP

- Bind Distinguished Name
  Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server.
- Bind Password
  Specifies the password to access the authentication server.
- Base Distinguished Name
  Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **`dc=ou,dc=domain,dc=com`**.

Others

- Bind Distinguished Name
  Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server that you have configured.
- Bind Password
  Specifies the password to access the authentication server.
- Base Distinguished Name
  Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **`dc=ou,dc=domain,dc=com`**.

- Protocol Version
  Specifies the Lightweight Directory Access Protocol (LDAP) version that is supported by your authentication server. You can specify whether the protocol version must be automatically detected or set the version to 2 or 3.
- User Name Attribute
  Specifies the name of the attribute in the authentication server that contains user login names to be authenticated by the management server.
- Group Membership Attribute
  Specifies a value that assigns the management server group membership to remote users based on an attribute and value specified in the user's authentication server.
- UGID
  If the remote users are included as members of a GroupOfUniqueNames object in the authentication server, this option enables you to assign the management server group membership to the remote users based on a specified attribute in that GroupOfUniqueNames object.
- Member
  Specifies the attribute name that your authentication server uses to store information about the individual members of a group.
- Group Object Class
  Specifies the object class of all groups in the remote authentication server.

Note: If you want to modify the authentication service, ensure that you delete any existing authentication servers and add new authentication servers.

Servers area

The Servers area displays the authentication servers that the management server communicates with to find and authenticate remote users. The credentials for remote users or groups are maintained by the authentication server.

Command buttons

Enables you to add, edit, or delete authentication servers.

- Add
  Enables you to add an authentication server.
  If the authentication server that you are adding is part of an high-availability pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.
- Edit
  Enables you to edit the settings for a selected authentication server.
- Delete
  Deletes the selected authentication servers.

Name or IP Address

Displays the host name or IP address of the authentication server that is used to authenticate the user on the management server.

Port

Displays the port number of the authentication server.

Test Authentication area

The Test Authentication area enables you to test your configuration.

Test

Validates the configuration of your authentication server by authenticating a remote user or group.

While testing, if you specify only the user name, the management server searches for the remote user in the authentication server, but does not authenticate the user. If you specify both the user name and password, the management server searches and authenticates the remote user.

You cannot test the authentication if remote authentication is disabled.

Command buttons

The command buttons enable you to save or cancel the setup options:

Restore to Factory Defaults

Enables you to restore the configuration settings to the factory default values.

Save

Saves the configuration settings for the selected option.

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Enabling remote authentication on page 395
Setting up authentication services on page 397
Adding authentication servers on page 398
Editing authentication servers on page 400
Testing the configuration of authentication servers on page 399
Deleting authentication servers on page 400

Managing security certificates

You can configure HTTPS in the Unified Manager server to monitor and manage your clusters over a secure connection.

Viewing the HTTPS security certificate

You can compare the HTTPS certificate details with the retrieved certificate in your browser to ensure that your browser's encrypted connection to Unified Manager is not being intercepted. You can also view the certificate to verify the content of a regenerated certificate, or to view alternate URL names from which you can access Unified Manager.

Before you begin

You must have the OnCommand Administrator role.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration > Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server > HTTPS**.
4. Click **View HTTPS Certificate**.

   To view detailed information about the security certificate, you can view the certificate in your browser.

Related tasks

Adding a user on page 379
Generating an HTTPS security certificate on page 406
Downloading an HTTPS certificate signing request on page 408
Installing an HTTPS security certificate on page 408
Restarting the Unified Manager virtual machine on page 405

Restarting the Unified Manager virtual machine

You can restart the Unified Manager virtual machine (VM) from the maintenance console. You must restart the VM after generating a new security certificate, or if there is a problem with the VM.

Before you begin

- The virtual appliance must be powered on.
- You must be logged in to the NetApp maintenance console as the maintenance user.

About this task

You can also restart the virtual machine from vSphere by using the VMware **Restart Guest** option.

Steps

1. In the maintenance console, select **System Configuration > Reboot Virtual Machine**.

2. Start the Unified Manager graphical user interface (GUI) from your browser, and log in.

Generating an HTTPS security certificate

You might generate a new HTTPS security certificate for multiple reasons, including when you want to sign with a different Certificate Authority or when the current security certificate has expired. The new certificate replaces the existing certificate.

Before you begin

You must have the OnCommand Administrator role.

About this task

> **Attention:** If connections that enable performance monitoring are currently configured between the Unified Manager server and one or more Performance Manager servers, executing this task invalidates those connections and deactivates any further performance monitoring updates from Performance Manager to the Unified Manager web UI. You must reactivate those connections after completing this task.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server** > **HTTPS**.
4. Click **Regenerate HTTPS Certificate**.

   Important: You must restart the Unified Manager virtual machine before the new certificate takes effect. You can use the **System Configuration** option in the NetApp maintenance console.

After you finish

After generating a new certificate, you can verify the new certificate information by viewing the HTTPS certificate.

If you need to reactivate performance monitoring updates from Performance Manager servers to the Unified Manager server, you must delete the connections that were invalidated by this task and reconfigure new connections.

Related tasks

Adding a user on page 379
Viewing the HTTPS security certificate on page 405
Downloading an HTTPS certificate signing request on page 408
Installing an HTTPS security certificate on page 408
Accessing the maintenance console using Secure Shell on page 407
Accessing the maintenance console using the vSphere VM console on page 407

Related references

Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring on page 465

Accessing the maintenance console using Secure Shell

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

Before you begin

You must have installed and configured Unified Manager.

You must be logged in as the maintenance user.

About this task

Note: You cannot perform maintenance console operations if Unified Manager is installed on Red Hat Enterprise Linux.

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

Steps

1. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance.
2. Log in to the maintenance console using your maintenance user name and password.

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

Accessing the maintenance console using the vSphere VM console

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

Before you begin

You must be the maintenance user. The virtual appliance must be powered on to access the maintenance console.

About this task

Note: You cannot perform maintenance console operations if Unified Manager is installed on Red Hat Enterprise Linux.

Steps

1. In vSphere Client, locate the Unified Manager virtual appliance.
2. Click the **Console** tab.
3. Click inside the console window to log in.
4. Log in to the maintenance console using your user name and password.

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

Downloading an HTTPS certificate signing request

You can download a certification request for the current HTTPS security certificate so that you can provide the file to a Certificate Authority (CA) to sign. A CA-signed certificate helps prevent man-in-the-middle attacks, and provides better security protection than a self-signed certificate does.

Before you begin

You must have the OnCommand Administrator role.

Steps

1. Click **Health** > **Administration** > **Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server** > **HTTPS**.
3. Click **Download HTTPS Certificate Signing Request**.
4. Save the `<hostname>.csr` file.

After you finish

You can provide the `<hostname>.csr` file to a Certificate Authority to sign, and then install the signed certificate.

Related tasks

Adding a user on page 379
Viewing the HTTPS security certificate on page 405
Generating an HTTPS security certificate on page 406
Installing an HTTPS security certificate on page 408

Installing an HTTPS security certificate

You can upload and install a security certificate after a Certificate Authority (CA) has signed and returned the security certificate. The file that you upload and install must be a signed version of the existing self-signed certificate. A CA-signed certificate helps prevent man-in-the middle attacks, and provides better security protection than a self-signed certificate.

Before you begin

- You must have downloaded the Certificate Signing Request file, and had it signed by a Certificate Authority.
- You must have saved the certificate chain in PEM format.
- You must have included all certificates in the chain, from the server certificate to the root signing certificate.
- You must have the OnCommand Administrator role.

About this task

Attention: If connections that enable performance monitoring are configured between the Unified Manager server and one or more Performance Manager servers, executing this task invalidates those connections, and deactivates any further performance monitoring updates from Performance

Manager servers to the Unified Manager web UI. You must reactivate those connections after completing this task.

Steps

1. Click ▦ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server** > **HTTPS**.
4. Click **Install HTTPS Certificate**.

Example certificate chain

The following example shows how the certificate chain file might appear:

```
-----BEGIN CERTIFICATE-----
<Server certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Root signing certificate>
-----END CERTIFICATE-----
```

After you finish

If you have to reactivate performance monitoring updates from Performance Manager servers to the Unified Manager server, you must delete the connections that were invalidated by this task, and then reconfigure new connections.

Related tasks

Adding a user on page 379
Viewing the HTTPS security certificate on page 405
Generating an HTTPS security certificate on page 406
Downloading an HTTPS certificate signing request on page 408

Related references

Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring on page 465

Description of security certificates windows and dialog boxes

You can use the Setup Options dialog box to view the current security certificates and to generate new HTTPS certificates.

HTTPS Setup Options dialog box

The HTTPS area in the Setup Options dialog box enables you to view the current security certificate, download a certificate signing request, to generate a new HTTPS certificate, or install a new HTTPS certificate.

You must have the OnCommand Administrator or Storage Administrator role.

HTTPS Certificate

You can perform the following operations:

View HTTPS Certificate

Enables you to view the current HTTPS certificate. If you have not generated a new HTTPS certificate, this is the certificate that was generated with your installation.

Regenerate HTTPS Certificate

Enables you to generate an HTTPS certificate, which replaces the previous security certificate. The new certificate is in effect after you restart the management server.

Download HTTPS Certificate Signing Request

Downloads a certification request for the currently installed HTTPS certificate. Your browser prompts you to save the `<hostname>.csr` file so that you can provide the file to a Certificate Authority to sign.

Install HTTPS Certificate

Enables you to upload and install a security certificate after a Certificate Authority has signed and returned it. The new certificate is in effect after you restart the management server.

Command buttons

The command buttons enable you to save or cancel the setup options:

Restore to Factory Defaults

Enables you to restore the configuration settings to the factory default values.

Save

Saves the configuration settings for the selected option.

Save and Close

Saves the configuration settings for the selected option and closes the Setup Options dialog box.

Cancel

Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Viewing the HTTPS security certificate on page 405
Generating an HTTPS security certificate on page 406
Downloading an HTTPS certificate signing request on page 408
Installing an HTTPS security certificate on page 408
Working with HTTPS security certificates on page 40

Managing reports

OnCommand Unified Manager enables you to create and manage reports so that you can view customized information about the capacity and utilization of storage objects and events related to storage objects.

Scheduling reports

You can schedule your reports from the Reports details page and email the scheduled reports to one or more recipients in a particular format at a specified frequency. For example, you can schedule a report to be sent as email, in the PDF format, every Monday.

Steps

1. From the **Reports details** page, click **Actions** > **Schedule Report**.
2. From the **Schedule Report** dialog box, you can select one of the preferred schedules for your report:

If you want to...	Then...
Select any schedule from the existing list of schedules	Use Existing Schedules
Create a new schedule	Enter the schedule name, specify the email address, select the report format and frequency in the specific fields. You can specify one or more email addresses, separated by commas. The *PDF* option is selected as the default report format. The *Hourly* option is selected as the default frequency.

3. Click **Schedule**.

Related concepts

What report scheduling is on page 416

Related references

Schedule Report dialog box on page 455

Sharing reports

You can email and share your reports with one or more users.

Steps

1. From the **Reports details** page, click **Actions** > **Share**.
2. In the **Share Report** dialog box, specify the email address of the recipient with whom you want to share the report.

   You can specify one or more email addresses, separated by commas.
3. Specify the subject of the email. By default, the name of the report appears as the subject of the email.
4. Select the report format.

The *PDF* option is selected as the default report format. If the XHTML format is selected, open the report that is sent by email by using a supported web browser.

5. Click **Share**.

Related concepts

What report sharing is on page 417

Related references

Share Report dialog box on page 456

Managing report schedules

You can manage your report schedules from Manage Report Schedules dialog box. You can add a new schedule and view, modify, or delete existing schedules.

Steps

1. In the **Reports** page, click **Manage Report Schedules**.
2. In the **Manage Report Schedules** dialog box,

If you want to...	Then...
View or modify existing schedule	**a.** Select the schedule from the list displayed in the left pane. **b.** The schedule details are displayed. **c.** Click **Save** or **Save and Close**.
Delete existing schedule	**a.** Select the schedule from the list displayed in the left pane. **b.** The schedule details are displayed. Make the necessary changes. **c.** Click **Delete Schedule**.
Add new schedule	**a.** Click **Add Schedule**. **b.** New schedule form appears in the right pane. **c.** Enter its specific details such as recipient schedule name, email address, report format, frequency and the reports. **d.** Click **Save**. The new schedule will be added in the Schedules list.

Related concepts

What report scheduling is on page 416

Related references

Manage Report Schedules dialog box on page 457

Customizing a report

You can customize reports in the Reports details page and then save the customized report with a different name.

Steps

1. Click **Reports**.
2. Select the type of report you want to customize, and then click **Run Report**.
3. Customize the report as necessary, and then click **Actions > Save Customized Report As**.
4. In the **Save Customized Report As** dialog box, enter a name for the customized report and a brief description about the customization.

   By default, the current report name is displayed.
5. Click **Save**.

   If you receive the error message "Failed to save the custom report. The required file was not created", wait a few moments, and then click **Save** again. This issue has been seen when there is a slow connection between the web browser and the Unified Manager server.

Result

The customized report is saved and displayed in its respective report category in the Reports details page.

Related concepts

What reports do on page 414

Related references

Save Customized Report As dialog box on page 458

Editing a customized report

You can make additional changes to an already customized report and save the report. You cannot change the name of the report after you have saved it.

Steps

1. Click **Reports**.
2. Select the type of report you want to customize and click **Run Report**.
3. To save the changes, click **Actions > Save Custom Report**.
4. In the **Save Custom Report** dialog box, enter a brief description about the changes made on the custom report and click **Save**.

Related concepts

What reports do on page 414

Related references

Save Custom Report dialog box on page 459

Importing reports

If you have created a report outside of Unified Manager, you can import and save the report file to use with Unified Manager.

Before you begin

You must have the OnCommand Administrator role.

Steps

1. From the **Reports** page, click **Import Report**.
2. In the **Import Report** dialog box, click **Browse** and select the file you want to import.
3. Click **Import**.

   If you cannot import the report, you can check the log file to find the error causing the issue.

Related concepts

What report importing is on page 417

Related references

Import Report dialog box on page 460

Understanding more about reports

You can use the option to run, delete, export, and import reports. You can also create custom reports and save the customized report. You can perform additional operations such as filtering, sorting, grouping, and formatting.

What reports do

Reports display detailed information about storage objects, which enable you to review and identify potential issues.

You can save, delete, share, schedule, and import reports. You can also search for specific reports. You can customize reports to address specific use cases, and save the customized report for future use. You can perform additional operations such as filtering, sorting, grouping, and formatting.

By default, each report group is displayed by report type and description. You can run reports to view a specific report group.

After you run a report, you can further customize it and save the customized report. You can view the custom reports that are saved in the Reports page, grouped under the specific report category.

You can schedule reports to be sent, or share reports in one of the supported formats: PDF, XHTML, CSV, XLS, or text.

You can export reports in different formats and save them on your desktop. You can export individual column data from the generated reports.

You can import report design files (`.rptdesign` files), and save the imported reports in the Reports page. You can delete custom and imported reports.

You can import the following reports:

- Reports with multiple headers that have a column span set to one
- Reports with charts only
- Reports with lists and grid only

Reports in text, CSV, and Excel formats are supported in the following scenarios:

- Table element only in the `.rptdesign` file
- A table with just one header as a row

You cannot import reports that have a column span of more than one. If a report in text, CSV, or Excel format has more than a one-header row, only the first header row is considered, and the remaining rows are ignored.

Database views

You can see the database views that are available in the ocum_report database by using the report import functionality. You can use the import report functionality to see the following database views:

- aggregate
- aggregatecapacityhistorymonthview
- aggregatecapacityhistoryweekview
- aggregatecapacityhistoryyearview
- annotation
- annotationofclusterview
- annotationofvolumeview
- annotationofvserverview
- cifsshare
- cifsshareacl
- cluster
- clusterlicensedetailview
- clusterlicenseview
- clusternode
- disk
- diskaggregatemapping
- exportpolicy
- exportrule
- fcpport
- lif
- lun
- networkport

- protectunprotectvolumeview
- qtree
- qtreequota
- storagepool
- storagepoolaggregatemapping
- svm
- svmaggregatemapping
- unprotectvolumeview
- userquota
- volume
- volumecapacityhistorymonthview
- volumecapacityhistoryweekview
- volumecapacityhistoryyearview
- volumeoutgoingrelationshipview
- volumeprotectionview
- volumerelationshiphistoryaggweekhourview
- volumerelationshiphistorychartcountview
- volumerelationshiphistorychartview
- volumerelationshiphistorydayview
- volumerelationshiphistorymonthview
- volumerelationshiphistoryview
- volumerelationshiphistoryweekview
- volumerelationshiphistoryyearview
- volumerelationshipinventoryview
- volumerelationships

Related references

Reports page on page 439

What report scheduling is

You can schedule a report to be generated at a specific date and time by using the **Schedule** option. The report is automatically sent by email to one or more recipients as per the schedule.

By scheduling a report, you can minimize the effort of generating and sending the reports manually. You can ensure that the current status of the storage is monitored at specified intervals by the administrators who are not otherwise notified by Unified Manager.

What report sharing is

You can share a report with one or more users through email using the **Share** option.

You must save the report prior to sharing it to ensure that the recent changes you made to the report is displayed.

You can share the report in any desired format. The **Share** option helps you to share reports through email instantly, even with persons who do not have access to Unified Manager but has a valid email address.

What report importing is

You can import a report using the **Import Report** option from Unified Manager and save the imported report with a name and a brief description. By importing reports, you can have additional reports other than the available reports in Unified Manager.

You can import a `.rptdesign` file that is already created. You can run, share, schedule, and delete an imported report.

Unified Manager stores the import report log files in jboss.log, ocum-report.log, and ocumserver-debug.log.

Note: Customer support will not assist with designing reports, but they will support you with issues faced during an import report.

The import report feature includes the following support:

- Reports with multiple headers, in which the column span is set to 1 (`colspan=1`)
- Reports with charts only
- Reports with lists and grid only

Reports in text, CSV, and Excel formats are supported in the following scenarios:

- Table element only in the `.rptdesign` file
- A table with only one header row

Note: You cannot import reports that have a column span of more than 1. If a report in text, CSV, or Excel format has more than one-header row, only the first header row is considered, and the rest are ignored.

Report customizations

You can customize various Unified Manager reports based on storage and utilization capacity, events, cluster inventory, NFS exports, or SVM inventory. .

Storage Summary report customizations

You can customize Storage Summary reports to view and analyze information about storage capacity in HA pairs. You can use filters to display storage utilization by cluster model, capacity of the most unassigned LUNs, and capacity of available HA pairs to provision new volumes and LUNs.

Customizing the Storage Summary report to view capacity by cluster models

You can customize the Storage Summary report to analyze storage capacity and utilization of clusters, and to view aggregates included in the total raw capacity.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster, perform the following steps:
    a. Click in the column that needs to be ungrouped.
    b. Click the ••• icon.
    c. Select **Group > Delete Inner Group**.
2. To group the report by the model name, perform the following steps:
    a. Click in the **Model** column and click the ••• icon.
    b. Select **Group > Add Group**.
3. To add aggregates to the total raw capacity, perform the following steps:
    a. Click in the **Total Raw Capacity** column and click the ••• icon.
    b. Select **Aggregation**.
    c. In the **Aggregation** dialog box, clear the **table level** check box and select the **group level** check box.
    d. Enter a label name in the **Enter Label** field, if required.
4. Click **OK**.
5. To add aggregates to the other columns in the report, repeat Steps 3 and 4.

Customizing the Storage Summary report to analyze cluster capacity based on Data ONTAP version

You can customize the Storage Summary report to group clusters by Data ONTAP version, and to view aggregates relating to your total raw capacity.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove grouping by cluster, perform the following steps:
    a. Click in the column that needs to be ungrouped.
    b. Click ••• (menu icon).
    c. Select **Group > Delete Inner Group option**.

2. To group the report by the Data ONTAP version, perform the following steps:
    a. Click in the **OS version** column and select the ⋯ icon.
    b. Select **Group > Add Group**.
3. To add aggregates to the total raw capacity, perform the following steps:
    a. Click in the **Total Raw Capacity** column and click the ⋯ icon.
    b. Select **Aggregation**.
    c. In the **Aggregation** dialog box, clear the **table level** check box and select the **group level** check box.
    d. Enter a label name in the **Enter Label** field, if required.
4. Click **OK**.

Customizing the Storage Summary report to analyze clusters with the most unallocated LUN capacity

You can customize the Storage Summary report to analyze the storage utilization of clusters, which enables you to locate the LUNs with the most unallocated capacity.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove grouping by cluster, perform the following steps:
    a. Click in the column that needs to be ungrouped.
    b. Click the ⋯ icon.
    c. Select **Group > Delete Inner Group**.
2. To sort HA pairs that have the most unallocated LUN capacity, click in the **Unallocated LUN Capacity (TB)** column, and click the ⋯ icon.
3. Select **Filter > Top/Bottom N**.
4. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
5. Click **OK**.

Customizing the Storage Summary report to analyze HA pairs for available capacity to provision new volume and LUNs

You can customize the Storage Summary report to display available HA pairs that have capacity, so that you can provision new volumes and LUNs. The report displays HA pairs sorted in order of decreasing aggregate unused capacity.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove grouping by cluster, perform the following steps:
   a. Click in the column that needs to be ungrouped.
   b. Click the ••• icon.
   c. Select **Group > Delete Inner Group**.
2. To sort HA pairs with available capacity, click in the **Aggregate Unused Capacity (TB)** column, and click the ••• icon.
3. Select **Filter > Top/Bottom N**.
4. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
5. Click **OK**.

Aggregate Capacity and Utilization Report customizations

You can customize reports to display a variety of information about aggregates.

Customizing the Aggregate Capacity and Utilization report to view aggregates reaching full capacity

You can customize the Aggregate Capacity and Utilization report to display aggregates sorted by increasing order of aggregate capacity utilization. This enables you to view the aggregates reaching full capacity.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove grouping by cluster and by HA pair, perform the following steps:
   a. Click in the columns that need to be ungrouped.
   b. Click the ••• icon.
   c. Select **Group > Delete Inner Group option**.
2. To sort the aggregates reaching full capacity, click in the **Days To Full** column, and click the ••• icon.
3. Select **Filter > Top/Bottom N**.
4. In the **Top/Bottom N** dialog box, select **Bottom N** from the **Filter** field and enter a value in the text field.
5. Click **OK**.

Customizing the Aggregate Capacity and Utilization report to display aggregates with the nearly full threshold breached

You can customize the Aggregate Capacity and Utilization report to display the top aggregates, sorted by decreasing order of Snapshot copy overflow percentage. This enables you to view the storage space still available in the aggregates.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster or HA pair, perform the following steps:
    a. Click in the column that needs to be ungrouped.
    b. Click the ⋯ icon.
    c. Select **Group > Delete Inner Group**.
2. To display the difference between the used data percentage and the nearly full threshold, add a new column:
    a. Select a column and click the ⋯ icon.
    b. Select **Column > New Computed Column**.
    c. In the **New Computed Column** dialog box, enter a column label.
    d. From the Select Category list, select **Math**.
    e. From the **Select Function** list, select **DIFFERENCE**.
    f. From the Column 1 list, select **Space Nearly Full Threshold (%)**.
    g. From the Column 2 list, select **Used Data%**.
    h. Click **OK**.
3. To filter the values greater than 0 in the new column, click in the **New computed column** and open the **Filter** dialog box by clicking the ▼ icon.
4. From the **Condition** drop-down list, select **Greater Than**.
5. In the **Value** field, type
   `0`
   and click **OK**.
6. To sort the values, click in the **New computed column** and click the ⋯ icon.
7. Select **Filter > Top/Bottom N**.
8. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
9. Click **OK**.

Customizing the Aggregate Capacity and Utilization report to display aggregates with overcommitted threshold breached

You can customize the Aggregate Capacity and Utilization report to display the aggregates sorted by overcommitted capacity percentage, which enables you to view the storage space still available in the aggregates.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster or HA pair, perform the following steps:
    a. Click in the column that needs to be ungrouped.
    b. Click the ⋯ icon.
    c. Select **Group > Delete Inner Group**.
2. To display the difference between the overcommitted used percentage and the overcommitted threshold, add a new column.
    a. Select a column and click ⋯.
    b. Select **Column > New Computed Column**.
    c. In the **New Computed Column** dialog box, enter a column label.
    d. From the Select Category list, select **Math**.
    e. From the **Select Function** list, select **DIFFERENCE**.
    f. From the Column 1 list, select **Overcommitted Threshold (%)**.
    g. From the Column 2 list, select **Overcommitted Capacity %**.
    h. Click **OK**.
3. To filter the values greater than zero in the new column, click in the **New computed column** and open the **Filter** dialog box by clicking the ▼ icon.
4. From the **Condition** list, select **Greater Than**.
5. In the **Value** field, type
   `0`
   and click **OK**.
6. To sort the values, click inside **New computed column** and click the ⋯ icon.
7. Select **Filter > Top/Bottom N**.
8. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
9. Click **OK**.

Customizing the Aggregate Capacity and Utilization report to display aggregates with non-compliant configuration compliance

You can customize the Aggregate Capacity and Utilization report to display the aggregates filtered by the full threshold. This enables you to view the aggregates that might not comply with company policies.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster or HA pair, perform the following steps:
    a. Click in the column that needs to be ungrouped.
    b. Click the ••• icon.
    c. Select **Group > Delete Inner Group**.
2. To filter aggregates threshold not exceeding 85%, click in the **Space Full Threshold** column and open the **Filter** dialog box by clicking the ▼ icon.
3. From the **Condition** list, select **Greater Than**.
4. Click **Select Values** and select **85**.
5. Click **OK**.

Volume Capacity and Utilization report customizations

You can create reports to monitor a variety of capacity and utilization information about volumes. For example, you can create reports to display volumes used, total capacity, daily growth rate, and Snapshot copy capacity, which can help you to determine if a volume is running out of space or whether it is being overutilized or underutilized.

Customizing the Volume Capacity and Utilization report to display volumes nearing full capacity with Snapshot Autodelete turned off

You can customize the Volume Capacity and Utilization report to display volumes sorted by increasing order of their volume capacity utilization. This enables you to display volumes reaching their full capacity.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by SVM, cluster, or volume, perform the following steps:
    a. Click in the column that needs to be ungrouped.
    b. Click the ••• icon.
    c. Select **Group > Delete Inner Group**.

2. To sort volumes that are nearing full capacity, click in the **Days To Full** column, and click the ⬆ icon.
3. To filter volumes that have Snapshot Autodelete turned off, click in the **Snapshot Autodelete** column and open the **Filter** dialog box by clicking the ▼ icon.
4. From the **Condition** list, select **Equal To**.
5. Click **Select Values** and select **Disabled**.
6. Click **OK**.

Customizing the Volume Capacity and Utilization report to display the least consumed volumes with thin provisioning disabled

You can customize the Volume Capacity and Utilization report to display volumes based on their volume consumption.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by SVM, cluster, or volume, perform the following steps:
   1. Click in the column that needs to be ungrouped.
   2. Click the ⋯ icon.
   3. Select **Group > Delete Inner Group**.
2. To sort volumes based on percentage consumed, click in the **Used Data %** column, and click the ⬆ icon.
3. To filter volumes with thin provisioning disabled, click in the **Thin Provisioned** column and open the **Filter** dialog box by clicking the ▼ icon.
4. From the **Condition** list, select **Equal To**.
5. Click **Select Values** and select **No**.
6. Click **OK**.

Customizing the Volume Capacity and Utilization report to display volumes with noncompliant configuration

You can customize the Volume Capacity and Utilization report to display volumes that are not compliant with company policies. For example, if you must have deduplication enabled on all volumes, you can create a report listing all volumes where deduplication is disabled.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by SVM, cluster, or volume, perform the following steps:
    a. Click in the column that needs to be ungrouped.
    b. Click the ••• icon.
    c. Select **Group > Delete Inner Group**.
2. Hide all columns except for the Cluster, Storage Virtual Machine, Volume, Deduplication, and Deduplication Space Savings (GB) columns:
    a. Click in the column and click the ••• icon.
    b. From the menu, select **Column > Hide Column**.
3. To filter volumes that deduplication disabled, click in the **Deduplication** column and open the **Filter** dialog box by clicking the ▼ icon.
4. From the **Condition** list, select **Equal To**.
5. Click **Select Values** and select **Disabled**.
6. Click **OK**.
7. To sort volumes based on deduplication space savings, click in the **Deduplication Space Savings (GB)** column and click the ⬇ icon.

Qtree Capacity and Utilization report customization

You can create customized reports to analyze capacity and utilization of the system's qtrees. For example, you can create reports to sort qtrees to determine whether any have breached the disk or file soft limit.

Customizing the Qtree Capacity and Utilization report to display qtrees that have breached the disk soft limit

You can customize the Qtree Capacity and Utilization report to display qtrees that have breached the disk soft limit. You can filter and sort by disk used, disk hard limit, and disk soft limit.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by SVM, cluster or volume, perform the following steps:
    a. Click in the columns that need to be ungrouped.
    b. Click the ••• icon.
    c. Select **Group > Delete Inner Group**.
2. To filter qtrees that do not have an unlimited disk hard limit, click in the **Disk Hard Limit** column and open the **Filter** dialog box by clicking the ▼ icon.
    a. From the **Condition** drop-down list, select **Not Equal To**.

   b. Click **Select Values** and select **Unlimited**.
   c. Click **Ok**.
3. To filter qtrees that do not have an unlimited disk soft limit, click in the **Disk Soft Limit** column and open the **Filter** dialog box by clicking the ▼ icon.
   a. From the **Condition** drop-down list, select **Not Equal To**.
   b. Click **Select Values** and select **Unlimited**.
   c. Click **Ok**.
4. To add a column for qtrees that have breached the disk soft limit, perform the following steps:
   a. Click in the **Disk Soft Limit** column, click the ⋯ icon, and select **Column** > **New Computed Column**
   b. In the **New Computed Column** dialog box, type `Breached Disk Soft Limit Capacity` in the **Column Label** field.
   c. From the Select Category list, select **Text**.
   d. From the **Select Function** drop-down list, select **Advanced**.
   e. In the **Enter Expression** field, type

   ```
   IF(([qtreeDiskUsedPercent] *[diskLimit]/100 > [softDiskLimit]), "Yes",
   "No")
   ```

   .
   f. Click **OK**.
5. To filter qtrees that have breached the soft disk limit, click in the **Breached Disk Soft Limit Capacity** column and open the **Filter** dialog box by clicking the ▼ icon.
   a. From the **Condition** drop-down list, select **Equal To**.
   b. Click **Select Values** and select **Yes**.
   c. Click **Ok**.

Customizing the Qtree Capacity and Utilization report to display qtrees that have breached the file soft limit

You can customize the Qtree Capacity and Utilization report to display qtrees that have breached the file soft limit. You can filter and sort by file used, file hard limit, and file soft limit.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by SVM, cluster or volume, perform the following steps:
   a. Click in the columns that need to be ungrouped.
   b. Click the ⋯ icon.
   c. Select **Group** > **Delete Inner Group**.

2. To filter qtrees that do not have an unlimited file hard limit, click in the **File Hard Limit** column and open the **Filter** dialog box by clicking the ▼ icon.
   a. From the **Condition** drop-down list, select **Not Equal To**.
   b. Click **Select Values** and select **Unlimited**.
   c. Click **Ok**.
3. To filter qtrees that do not have an unlimited file soft limit, click in the **File Soft Limit** column and open the **Filter** dialog box by clicking the ▼ icon.
   a. From the **Condition** drop-down list, select **Not Equal To**.
   b. Click **Select Values** and select **Unlimited**.
   c. Click **Ok**.
4. To add a column for qtrees that have breached the file soft limit, perform the following steps:
   a. Click in the **File Soft Limit** column, click the ••• icon, and select **Column > New Computed Column**
   b. In the **New Computed Column** dialog box, type `Breached File Soft Limit Capacity` in the **Column Label** field.
   c. From the Select Category list, select **Text**.
   d. From the **Select Function** drop-down list, select **Advanced**.
   e. In the **Enter Expression** field, type

      ```
      IF(([qtreeFileUsedPercent]*[fileLimit]/100 > [softFileLimit]), "Yes",
      "No")
      ```
      .
   f. Click **OK**.
5. To filter qtrees that have breached the soft file limit, click in the **Breached File Soft Limit Capacity** column and open the **Filter** dialog box by clicking the ▼ icon.
   a. From the **Condition** drop-down list, select **Equal To**.
   b. Click **Select Values** and select **Yes**.
   c. Click **Ok**.

Events report customization

You can create reports to monitor outstanding events on a cluster.

Customizing the Events report to display events with a critical severity type

You can customize the Events report to display events filtered by their severity type, and by the events that have been unresolved for the longest period of time.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To filter events with critical severity type, click in the **Status** column and open the **Filter** dialog box by clicking the ▼ icon.
2. From the **Condition** list, select **Equal To**.
3. Click **Select Values** and select **Critical**.
4. Click **OK**.
5. To sort the events that are unresolved for the longest period of time, click in the **Days Outstanding** column, and click the ••• icon.
6. Select **Filter > Top/Bottom N**.
7. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
8. Click **OK**.

Customizing the Events report to display events on mission-critical objects

You can customize the Events report to display events filtered by mission-critical data priority.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To filter events with mission-critical data priority, click in the **Data Priority** column and open the **Filter** dialog box by clicking the ▼ icon.
2. From the **Condition** list, select **Equal To**.
3. Click **Select Values** and select **Mission-Critical**.
4. Click **OK**.

Customizing the Events report to display the top most discussed events

You can customize the Events report to display events that are most discussed.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To sort the events that are discussed the most, click in the **Notes** column and click the ••• icon.
2. Select **Filter > Top/Bottom N**.
3. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
4. Click **OK**.

Customizing the Events report to display incident events assigned to the admin

You can customize the Events report to display incident events that are assigned to the admin, filtered by the impact level and the admin name.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To filter incident events, click in the **Impact Level** column and open the **Filter** dialog box by clicking the icon.
2. From the **Condition** list, select **Equal To**.
3. Click **Select Values** and select **Incident**.
4. Click **OK**.
5. To assign these incidents to the admin, click in the **Assigned To** column and open the **Filter** dialog box by clicking the icon.
6. From the **Condition** drop-down list, select **Equal To**.
7. Click **Select Values** and select **Admin Name**.
8. Click **OK**.

Customizing the Events report to display events impacting availability

You can customize the Events report to display events that are categorized by the most incidents and are assigned to the admin. You can filter the report by the impact level and the admin name.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To filter availability events, click in the **Impact Area** column and open the **Filter** dialog box by clicking the icon.
2. From the **Condition** drop-down list, select **Equal To**.
3. Click **Select Values** and select **Incident**.
4. Click **OK**.

Customizing the Events report to display the top most acknowledged unresolved events

You can customize the Events report to display the most acknowledged events, filtered by the event state. You can sort them in decreasing order to display the number of outstanding days.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To filter acknowledged events, click in the **State** column and open the **Filter** dialog box by clicking the ▼ icon.
2. From the **Condition** drop-down list, select **Equal To**.
3. Click **Select Values** and select **Acknowledged**.
4. Click **OK**.
5. To further filter the report, click in the **Acknowledged By** column and open the **Filter** dialog box by clicking the ▼ icon.
6. From the **Condition** drop-down list, select **Equal To**.
7. Click **Select Values** and select **Name**.
8. Click **OK**.
9. To sort the events that are outstanding for the most number of days, click in the **Days Outstanding** column and click ⋯.
10. Select **Filter > Top/Bottom N**.
11. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
12. Click **OK**.

Cluster Inventory Report customizations

You can customize inventory reports to monitor for insufficient resources on clusters components. For example, you can customize reports to monitor information such as clusters that are nearing the SVM count limit, nodes that are running older versions of Data ONTAP, and nodes that are reaching the maximum disk limit.

Customizing the Cluster Inventory report to display clusters reaching SVM count limit

You can customize the Cluster Inventory report to display clusters, sorted by decreasing order of their SVM count.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster or node, perform the following steps:
   a. Click in the column that needs to be ungrouped.
   b. Click the ••• icon.
   c. Select **Group > Delete Inner Group**.
2. To sort clusters by SVM count, click in the **SVM Count** column and click •••.
   a. Click in the column that needs to be ungrouped.
   b. Click the ••• icon.
   c. Select **Group > Delete Inner Group option**.
3. Select **Filter > Top/Bottom N**.
4. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
5. Click **OK**.

Customizing the Cluster Inventory report to display nodes running older versions of Data ONTAP

You can customize the Cluster Inventory report to display nodes filtered by older Data ONTAP versions.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster, or node, perform the following steps:
   a. Click in the column that needs to be ungrouped.
   b. Click the ••• icon.
   c. Select **Group > Delete Inner Group**.
2. To filter nodes not running Data ONTAP 8.3, click the **Data ONTAP version** column and open the **Filter** dialog box by clicking the ▼ icon.
3. From the **Condition** drop-down list, select **Not Equal To**.
4. Click **Select Values** and select **8.3**.
5. Click **OK**.

Customizing the Cluster Inventory report to display nodes reaching the maximum disk limit

You can customize the Cluster Inventory report to display a list of nodes that are reaching the maximum disk limit and sorted by increasing order.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster, or node, perform the following steps:
    a. Click in the columns that needs to be ungrouped.
    b. Click the ⋯ icon.
    c. Select **Group > Delete Inner Group**.
2. To move the **Disk Count** column next to the **Model** column, perform the following steps:
    a. Click in the **Disk Count** column.
    b. Click the ⋯ icon and select **Column > Reorder Columns**.
    c. In the **Reorder Columns** dialog box, use the **up** and **down** arrow keys to move the column to the required position.
3. To add a new computed column, perform the following steps:
    a. Select a column, click ⋯, and select **Column > New Computed Column**.
    b. In the **New Computed Column** dialog box, type `Maximum Disk Limit` in the **Column Label** field.
    c. From the Select Category list, select **Comparison**.
    d. From the **Select Function** list, select **Advanced**.
    e. In the **Enter Expression** field, type `IF([model]="FAS3250" , 960, 0)`.
    f. Click **OK**.
4. To add a second new column, perform the following steps:
    a. Select the **Maximum Disk Limit** column, click the ⋯ icon, and select **Column > New Computed Column**.
    b. In the **New Computed Column** dialog box, type `Available Volume` in the **Column Label** field.
    c. From the Select Category list, select **Math**.
    d. From the **Select Function** list, select **DIFFERENCE**.
    e. From the Column 1 list, select **Maximum Disk Limit**.
    f. From the Column 2 list, select **Disk Count**.
    g. Click **OK**.

5. To sort the values, click in the **Available Volume** column, and click the ⋯ icon.
6. Select **Filter > Top/Bottom N**.
7. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
8. Click **OK**.

NFS Export report customizations

You can customize NFS export reports to analyze information about NFS export policies and rules for volumes on your storage systems. For example, you can customize reports to display volumes with inaccessible junction paths and volumes with the default export policy.

Customizing the NFS Exports report to display a list of volumes that have an inaccessible junction path

You can customize the NFS Exports report to display a list of volumes that have an inaccessible junction path.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster or volume, perform the following steps:
   a. Click in the columns that need to be ungrouped.
   b. Click the ⋯ icon.
   c. Select **Group > Delete Inner Group**.
2. To filter volumes that have an inaccessible junction path, click in the **Junction Path Active** column and open the **Filter** dialog box by clicking the ▼ icon.
3. From the **Condition** list, select **Equal To**.
4. Click **Select Values** and select **No**.
5. Click **OK**.

Customizing the NFS Exports report to display a list of volumes with default export policy

You can customize the NFS Exports report to display a list of volumes with default export policy.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster or volume, perform the following steps:
   a. Click in the columns that need to be ungrouped.
   b. Click the ⋯ icon.

   c. Select **Group** > **Delete Inner Group**.

2. To filter volumes with default export policy, click the **Export Policy** column and open the **Filter** dialog box by clicking the ▼ icon.
3. From the **Condition** list, select **Equal To**.
4. Click **Select Values** and select **Default**.
5. Click **OK**.

SVM Inventory report customization

You can create SVM inventory reports to analyze volume information and to view overall health and storage availability. For example, you can create reports to display SVMs reaching the maximum volume count and to analyze stopped SVMs.

Customizing the SVM Inventory report to display a list of SVMs reaching maximum volume limit

You can customize the SVM Inventory report to display a list of SVMs that are reaching the maximum volume limit by sorting the volumes in increasing order.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To remove the grouping by cluster, perform the following steps:
   a. Click inside the column that needs to be ungrouped.
   b. Click the ••• icon.
   c. Select **Group** > **Delete Inner Group**.
2. To filter Storage Virtual Machines (SVMs) that do not have unlimited allowed volumes, click the **Maximum Allowed Volumes** column and open the **Filter** dialog box by clicking the ▼ icon.
3. In the **Data type** field, select **String** and click **OK**.
4. From the **Condition** drop-down list, select **Not Equal To**.
5. Click **Select Values** and select **Unlimited**.
6. To add a new computed column, perform the following steps:
   a. Select a column, click the ••• icon, and select **Column > New Computed Column**.
   b. In the **New Computed Column** dialog box, type **`Available Volume`** in the **Column Label** field.
   c. From the Select Category list, select **Math**.
   d. From the **Select Function** drop-down list, select **Advanced**.
   e. In the **Enter Expression** field, type
      `[maximumVolumes]-[volumeCount]`

    .

    f. Click **OK**.

7. To sort Storage Virtual Machines (SVMs) in ascending order, click in the **Available Volume** column, and click the ••• icon.
8. Select **Filter** > **Top/Bottom N**.
9. In the **Top/Bottom N** dialog box, select **Bottom N** from the **Filter** field and enter a value in the text field.
10. Click **OK**.

Customizing the SVM Inventory report to display a list of stopped Storage Virtual Machines (SVMs)

You can customize the SVM Inventory report to display a list of stopped Storage Virtual Machines (SVMs). The report filters the Storage Virtual Machines (SVMs) by their status.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To filter Storage Virtual Machines (SVMs) by status, click the **State** column and open the **Filter** dialog box by clicking the ▼ icon.
2. From the **Condition** list, select **Equal To**.
3. Click **Select Values** and select **Stopped**.
4. Click **OK**.

Volume Relationships Inventory report customizations

You can customize the Volume Relationships Inventory report to view the volume details that are filtered based on the source of failure. You can use filters to display volume relationships inventory details based on schedules, and to group volume inventory details based on issues.

Customizing the Volume Relationships Inventory report to view volumes grouped by source of failure

You can customize the Volume Relationships Inventory report to view volumes grouped by source of failure.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. Select the **Relationship Health** column.
2. To view the volume details for bad volumes, click the – sign next to the **Bad** column.
3. To view the volume details for good volumes, click the – sign next to the **Good** column.

Customizing the Volume Relationships Inventory report to view volumes grouped by issue

You can customize the Volume Relationships Inventory report to view volumes that are grouped according to the volume relationship health status.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. To filter volumes according to the volume relationship health status, select the **Relationship Health** column, and click the ▼ icon.
2. In the **Filter** dialog box, click **Select Values**, and then select the required value from the drop-down list.

   The volume details for the selected value are displayed.

Customizing the Volume Transfer Status report to view volumes based on schedules

You can customize the Volume Transfer Status report to view the volume details that are sorted based on different schedules. You can view, modify, or delete existing report schedules, and add new schedules for your reports.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. On the **Volume Transfer Status** report page, click **Manage Report Schedules**.
2. In the **Manage Report Schedules** dialog box, enter specific details such as recipient schedule name, email address, report format, frequency, and the reports.
3. Select **Inventory** as the Report Category.
4. Click **Save and Close**.

   The Volume Transfer Status report is automatically sent by email to one or more recipients as per the schedule.

Volume Transfer Status report customizations

You can customize the Volume Transfer Status report to view and analyze information about volume transfers at specific time intervals. You can use filters to view volume transfer details between two dates.

Customizing the Volume Transfer Status report to view volumes at specific time intervals

You can customize the Volume Transfer Status report to view the volume details at specific time intervals.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. Remove grouping by cluster:
    a. Click in the column that you want to ungroup.
    b. Click the ••• icon.
    c. Select **Group > Delete Inner Group**.
2. To view the volume details at a specific time interval, click in the **Start time** column, and then click the ▼ icon.
3. In the **Filter** dialog box, click **Select Values**, and then select the specific date and time from the drop-down list.

    The volume details for the selected time range are displayed.

Customizing the Volume Transfer Status report to view volumes grouped by time of occurrence

You can customize the Volume Transfer Status report to display the list of volumes grouped by time of occurrence between two dates.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. Remove grouping by cluster:
    a. In the column that has to be ungrouped, click the ••• icon.
    b. Select **Group > Delete Inner Group**.
2. In the **Start time** column, open the **Filter** dialog box by clicking the ▼ icon.
3. From the **Condition** drop-down list, select **Between**.
4. Click **Select Values**, and choose the **Date From** and **Date To** values.
5. Click **OK**.

Customizing the Volume Transfer Status report to view failed or successful volume transfers

You can customize the Volume Transfer Status report to view the details of failed or successful volume transfers.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. Remove grouping by cluster:
    a. Select the column that you want to ungroup.

   b. Click the ••• icon.

   c. Select **Group > Delete Inner Group**.

2. To sort the volume transfers according to failure or success, click in the **Operational Result** column, and then click the ••• icon.
3. Select **Filter**.
4. In the **Filter** dialog box, click **Select Values**, and then select either **Success** or **Failure**.

Volume Transfer Rate report customizations

You can customize the Volume Transfer Rate report to view the volume transfer details based on the total transfer size of the volume. You can also view the volume transfers for a specific day of the week.

Customizing the Volume Transfer Rate report to view volume transfers based on transfer size

You can customize the Volume Transfer Rate report to view the volume transfer details according to the total transfer size of the volume.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. Remove grouping by cluster:
   a. Select the column that you want to ungroup.
   b. Click the ••• icon.
   c. Select **Group > Delete Inner Group**.
2. To sort the volume transfers according to volume transfer size, click the **Total Transfer Size (GB)** column.

Customizing the Volume Transfer Rate report to view volume transfers grouped by day

You can customize the Volume Transfer Rate report to view the volume transfer details that are sorted by day.

About this task

You can also perform this task by going to the Reports page and selecting **Run Report** for the appropriate report.

Steps

1. Remove grouping by cluster:
   a. Select the column thatyou want to ungroup.
   b. Click the ••• icon.
   c. Select **Group > Delete Inner Group**.
2. To view the volume transfers for a specific day, click the **Day** column.

Description of report windows and dialog boxes

You can use the options to schedule, share, manage, save, and import the reports.

Reports page

The Reports page enables you to view detailed information about the reports that you generate. You can search for a specific report, save a report, and delete a report. You can also schedule, share, and import a report.

The Reports page displays categorized groups of reports about which you can obtain specific report details. By default, the report groups expand to display the report types, a report overview, and links that enable you to customize reports. Only one report can be viewed at a time. You can click the **Run Report** button to view a report for a specific group of reports.

The following is a list of report groups and report types that are displayed:

- Capacity Utilization Reports
    - Storage Summary
    - Aggregates Capacity and Utilization
    - Volumes Capacity and Utilization
    - Qtree Capacity and Utilization
- Operational Reports - Events
- Inventory Reports
    - Cluster Inventory
    - NFS Exports
    - SVM Inventory
- Imported Reports
- Data Protection Reports
    - Volume Data Protection Configuration
    - Volume Relationships Inventory
    - Volume Transfer Status
    - Volume Transfer Rate

Related references

Storage Summary report on page 440
Aggregate Capacity and Utilization report on page 441
Volume Capacity and Utilization report on page 443
Qtree Capacity and Utilization report on page 445
Events report on page 446
Cluster Inventory report on page 448
NFS Exports report on page 449
SVM Inventory report on page 451

Storage Summary report

The Storage Summary report enables you to view summarized information about storage capacity in the HA pairs. This information helps you to understand possible capacity risks and to take appropriate action to rebalance workload. Single-node cluster information is not visible in the report.

The Storage Summary report is displayed in two formats:

- Storage Summary report Chart view
- Storage Summary report tabular view

Storage Summary report Chart view

The chart shows the capacity trend of used and unused data capacity of the aggregates over a period of time. Total data capacity is displayed on the vertical (y) axis and the cluster name on the horizontal (x) axis. Therefore, each bar in the chart represents one cluster. You can view the details for specific points on the graph by positioning your cursor over a particular point.

Storage Summary report tabular view

Cluster Name

Displays the cluster name.

HA Pair

Displays the HA pair value obtained by forming two nodes.

Model

Displays the name of the model.

OS Version

Displays the version of Data ONTAP used.

Total Raw Capacity (TB)

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

Unconfigured Raw Capacity (TB)

Displays the unconfigured capacity of disks whose container type is other than aggregate, broken, spare, or shared. This capacity is always higher than the physical capacity of the disk in Data ONTAP. For example, consider a 2 TB disk. The physical capacity of the disk is 1.6 TB in Data ONTAP whereas the unconfigured raw capacity in OnCommand Unified Manager is 1.8 TB.

Aggregate Total Capacity (TB)

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

Aggregate Used Capacity (TB)

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

Aggregate Unused Capacity (TB)

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

Allocated LUN Capacity (TB)

Displays the capacity of LUNs that are mapped.

Unallocated LUN Capacity (TB)

Displays the capacity of all LUNs not mapped to the Host.

Volume Capacity (TB)

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

Volume Used Capacity (TB)

Displays the used capacity of the volumes.

Volume Unused Capacity (TB)

Displays the unused capacity of the volumes.

Volume Protection Capacity (TB)

Displays the capacity of volumes that have SnapMirror and SnapVault enabled.

Related concepts

Managing and monitoring clusters and cluster objects on page 144

Aggregate Capacity and Utilization report

The Aggregate Capacity and Utilization report enables you to view information about the capacity and utilization of aggregates in a cluster. This information enables you to understand possible capacity risks and also to view the configured, used, and unused capacity of aggregates.

Aggregate Capacity and Utilization report tabular view

Cluster

Displays the cluster name.

HA Pair

Displays the HA pair value obtained by forming two nodes.

Aggregate

Displays the aggregate name.

Total Data Capacity (GB)

Displays the total data capacity (used plus available).

Used Data Capacity (GB)

Displays the used data capacity.

Used Data %

Displays the used data capacity as a percentage.

Available Data Capacity (GB)

Displays the available data capacity.

Available Data %

Displays the available data capacity as a percentage.

Used Capacity Trend

Displays a chart of used data capacity of the aggregate over a period of time.

Daily Growth Rate %

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

Days To Full

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

Space Full Threshold

Displays the percentage at which an aggregate is full.

Space Nearly Full Threshold

Displays the percentage at which an aggregate is nearly full.

Growth Rate Threshold

Specifies the aggregate's growth rate is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event.

Growth Rate Sensitivity Threshold

Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated.

Days Until Full Threshold

Specifies the number of days remaining before the aggregate reaches full capacity.

Snapshot Reserve Total Capacity (GB)

Displays the total snapshot reserve capacity of the aggregate.

Snapshot Reserve Used Capacity (GB)

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

Snapshot Reserve Used %

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

Snapshot Reserve Available Capacity (GB)

Displays the amount of space available for Snapshot copies.

Snapshot Reserve Available %

Displays the amount of space available for Snapshot copies as a percentage.

Snapshot Copies Reserve Full Threshold

Specifies the percentage at which an aggregate has consumed all its space reserved for Snapshot copies.

Overcommitted Capacity %

Displays the aggregate overcommitment as a percentage.

Overcommitted Threshold %

Displays the percentage at which an aggregate is overcommitted.

Nearly Overcommitted Threshold %

Displays the percentage at which an aggregate is nearly overcommitted.

Type

Indicates whether the aggregate is a Flash Pool aggregate (combines HDDs and SSDs), or whether the disks in the aggregate are standard disks (HDDs only) or SSD disks (SSDs only).

For standard disks and SSD disks, this column is blank when the monitored storage system is running clustered Data ONTAP version earlier than 8.3.

RAID Type

Displays the RAID configuration type.

Aggregate State

Displays the current state of the aggregate.

Related references

Aggregates page on page 200

Volume Capacity and Utilization report

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

Volume Capacity and Utilization report tabular view

Cluster

Displays the cluster name.

Storage Virtual Machine

Displays the name of the Storage Virtual Machine (SVM) that contains the volume.

Volume

Displays the volume name.

Total Data Capacity (GB)

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

Used Data Capacity (GB)

Displays the used data capacity in a volume.

Used Data %

Displays the used data in a volume as a percentage.

Available Data Capacity (GB)

Displays the available data capacity in a volume.

Available Data %

Displays the available data capacity in a volume as a percentage.

Used Capacity Trend

Displays a chart of used data capacity of the aggregate over a period of time.

Daily Growth Rate %

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

Days To Full

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

Space Full Threshold %

Specifies the limit to the volume that is considered full.

Space Nearly Full Threshold %

Specifies the limit to the volume that is considered nearly full.

Growth Rate Threshold %

Specifies the aggregate's growth rate is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event.

Growth Rate Sensitivity Threshold

Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated.

Days Until Full Threshold

Specifies the number of days remaining before reaching full capacity.

Snapshot Overflow %

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

Snapshot Reserve Used Capacity (GB)

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

Snapshot Reserve Used %

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

Snapshot Reserve Available Capacity (GB)

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

Snapshot Reserve Available %

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

Snapshot Reserve Total Capacity (GB)

Displays the total Snapshot copy capacity in the volume.

Snapshot Copies Reserve Full Threshold %

Specifies the percentage at which the space reserved for Snapshot copies is considered full.

Snapshot Copies Count Threshold

Specifies the number of Snapshot copies that can be created on a volume before the system generates the Too Many Snapshot Copies event.

Snapshot Copies Days Until Full Threshold

Specifies the number of days remaining before the space reserved for Snapshot copies reaches full capacity.

Number Of Inodes

Displays the number of inodes in the volume.

Inode Utilization

Specifies the inode space used in the volume.

Inodes Full Threshold

Specifies the percentage at which a volume is considered to have consumed all of its inodes.

Inodes Nearly Full Threshold

Specifies the percentage at which a volume is considered to have consumed most of its inodes.

Quota Committed Capacity (GB)

Displays the space reserved in the volumes.

Quota Overcommitted Capacity (GB)

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

Quota Overcommitted Threshold %

Specifies the percentage at which the volume is nearly overcommitted.

Quota Nearly Overcommitted Threshold %

Specifies the percentage at which the volume space is nearly overcommitted.

Snapshot Autodelete

Displays whether automatic deletion of Snapshot copies is enabled or disabled.

Deduplication

Displays whether deduplication is enabled or disabled for the volume.

Deduplication Space Savings (GB)

Displays the amount of space saved in a volume by using deduplication.

Compression

Displays whether compression is enabled or disabled for the volume.

Compression Space Savings (GB)

Displays the amount of space saved in a volume by using compression.

Caching Policy

Displays the caching policy that is associated with the selected volume. The policy provides information about how flash pool caching occurs for the volume. See Volumes page for more information on caching policies.

Thin Provisioned

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

Autogrow

Displays whether the FlexVol volume automatically grows in size when it is out of space.

Space Guarantee

Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate.

State

Displays the state of the volume that is being exported.

Related references

Volumes page on page 211

Volume Thresholds Setup Options dialog box on page 132

Qtree Capacity and Utilization report

The Qtree Capacity and Utilization report enables you to analyze capacity and utilization of the system's qtrees to understand possible risks that might occur due to reduced cluster capacity.

Qtree Capacity and Utilization report tabular view

Cluster

Displays the name of the cluster containing the qtree.

Storage Virtual Machine

Displays the SVM name containing the qtree.

Volume

Displays the name of the volume containing the qtree.

Qtree

Displays the name of the qtree.

Quota type

Specifies if the quota is for a user, user group or a qtree.

User or Group

Displays the name of the user or user group. There will be multiple rows for each user and user group. When the quota type is qtree, then *Not Applicable* is displayed. If the quota is not set, then the column is empty.

Disk Used %

Displays the percentage of the disk space used. If a disk hard limit is set, this value is based on the disk hard limit. If the quota is set without a disk hard limit, the value is based on the volume data space. If the quota is not set or if the quotas are off on the volume to which the qtree belongs, then *Not applicable* is displayed.

Disk Hard Limit

Displays the maximum disk space allocated for the qtree. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as *Unlimited* if the quota is set without a disk hard limit, If the quota is not set, or if the quotas are off on the volume to which the qtree belongs.

Disk Soft Limit

Displays the disk space allocated for the qtree before a warning event is generated. The value is displayed as *Unlimited* if the quota is set without a disk soft limit, if the quota is not set, or if the quotas are off on the volume to which the qtree belongs.

Files Used %

Displays the percentage of files used in the qtree. If the file hard limit is set, this value is based on the file hard limit. The value is displayed as *Not applicable* if the quota is not set, or if the quota is set without a file hard limit, or if the quotas are off on the volume to which qtree belongs.

File Hard Limit

Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as *Unlimited* if the quota is set without a file hard limit, if the quota is not set, or if the quotas are off on the volume to which the qtree belongs.

File Soft Limit

Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as *Unlimited* if the quota is set without a file soft limit, if the quota is not set, or if the quotas are off on the volume to which the qtree belongs.

Related concepts

Managing quotas on page 136

Events report

The Events report enables you to view information about event trends over a specific time period. This information enables you to compare recent activity with any past operational activity, such as configuration changes, upgrades, and so on. The information also helps you to determine any outstanding events.

The Events report is displayed in two formats:

- Events report Chart view
- Events report tabular view

Events report Chart view

The Events Chart is displayed in two formats:

- Events Severity Trend
- Event Status Trending Per Day

Events Severity Trend

The chart shows the event severity trends for all open events over a time period. A count of events is displayed on the vertical (y) axis and the date is displayed on the horizontal (x) axis. You can view the details for specific points on the graph by positioning your cursor over a particular point. The details display the event severity, number of events of the specific severity type, and the date of the event.

The event severity types displayed are Critical, Error, and Warning. The event severities are differentiated by different colors. There can be the same number of events on the same date in different states.

Count

Displays a count of events.

Date

Displays the date. The x axis shows data from the time that the event occurred up to the present date. You can click and zoom the chart to get details.

Event Status Trending Per Day

The chart shows the event status trending per day over a period of time. A count of events is displayed on the vertical (y) axis and the date is displayed on the horizontal (x) axis. The details display the event state, number of events of the specific state, and the date of the event.

The event status are New, Acknowledged, and Resolved. The event status are differentiated by different colors.

The chart shows the new events generated daily on a cumulative basis in a bar graph represented in green color. The number of Acknowledged and Resolved events are shown as and when they are acknowledged and resolved on a daily basis.

There is a zoom functionality provided within the charts. You can use this feature to zoom a particular point in the chart for more clarity.

Events report tabular view

Source

Displays the source of an event.

Status

Displays the severity of the event. You can filter this column to display events of a specific severity type. The event severity types are Critical, Error, or Warning.

State

Displays the event state: New, Acknowledged, Resolved, or Obsolete. You can filter this column to show events of a specific state.

Event

Displays the event names.

Triggered Time

Displays the time when the event was generated. Both the time and the date are displayed.

Days Outstanding

Displays the number of days between an event occurring and its resolution or designation as Obsolete.

Source Type

Displays the object type (for example, Storage Virtual Machine (SVM), volume, or qtree) with which the event is associated.

Data Priority

Displays the annotation type, based on the priority of data of the storage object.

Impact Level

Displays whether the event is categorized as an incident, a risk, or information.

Impact Area

Displays whether the event is a capacity, availability, performance, protection, or configuration event.

Assigned To

Displays the name of the user to whom the event is assigned.

Assigned Time

Displays the time when the event was assigned to a user.

Notes

Displays the number of notes that are added for an event.

Acknowledged By

Displays the name of the user who acknowledged the event. The field is blank if the event is not acknowledged.

Acknowledged Time

Displays the time that has elapsed since the event was acknowledged. If the time elapsed exceeds a week, the timestamp displays when the event was acknowledged.

Resolved By

Displays the name of the user who resolved the event. The field is blank if the event is not resolved.

Resolved Time

Displays the time that has elapsed since the event was resolved. If the time elapsed exceeds a week, the timestamp displays when the event was resolved.

Obsoleted Time

Displays the time when the state of the event became Obsolete.

Related concepts

What events are on page 73

Related references

Events page on page 109

Cluster Inventory report

Cluster Inventory report provides information about available resources for cluster components for the purpose of understanding possible risks caused by insufficient resources.

Cluster Inventory report tabular view

Cluster

Displays the name of the cluster.

HA pair

Displays the HA pair value obtained by forming two nodes.

Node

Displays the name of the nodes.

Model

Displays the name of the model.

Data ONTAP version

Displays the version of Data ONTAP used.

All Flash Optimized

Displays whether node is configured to support only solid-state drives (SSDs).

Serial Number

Displays the serial number of the node.

Firmware Version

Displays the firmware version of the node.

SVM Count

Displays the number of SVMs contained by the cluster.

FC Port Count

Displays the number of FC ports contained by the node.

FCoE Port Count

Displays the number of FCoE ports contained by the node.

Ethernet Port Count

Displays the number of ethernet ports contained by the node.

Flash Card Count

Displays the number of flash cards installed on nodes in your data center so that you can monitor for potential problems.

Flash Card Size (GB)

Displays the size of the flash cards installed on nodes.

Disk Shelves Count

Displays the number of disk shelves contained by the node.

Disk Count

Displays the number of disks in a node.

Related concepts

Understanding clusters and cluster objects on page 153

Related references

Clusters page on page 168

NFS Exports report

NFS Exports report enables you to audit information about NFS export policies and its associated rules for volumes in your storage system.

NFS Exports report tabular view

Cluster

Displays the name of the cluster.

SVM

Displays the name of the SVM with NFS export policies.

Volume

Displays the name of the volume with NFS export policies.

Qtree

Displays the name of the qtree on a volume with NFS export policies.

Volume State

Displays the current state of the volume. The state can be Offline, Online, or Restricted.

- Offline
  Read or write access to the volume is not allowed.
- Online
  Read and write access to the volume is allowed.
- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Junction Path

Displays the path on which the volume is mounted.

Junction Path Active

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

Export policy

Displays the rules that define the access permission for volumes that are exported.

Rule Index

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

Access Protocols

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

Client Match

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

Read Only Access

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

Read Write Access

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

Security Style

Displays the access permission for volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

- UNIX (NFS clients)
  Files and directories in the volume have UNIX permissions.
- Unified
  Files and directories in the volume have a unified security style.
- NTFS (CIFS clients)
  Files and directories in the volume have Windows NTFS permissions.
- Mixed
  Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

Unix Permission

Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.

Related references

Storage Virtual Machine details page on page 185

SVM Inventory report

SVM Inventory report enables you to analyze SVM volume configuration limits and overall health to understand risks to future storage availability.

SVM Inventory report tabular view

Cluster

Displays the name of the cluster containing the SVM.

Storage Virtual Machine

Displays the name of the SVM.

State

Displays the current administrative state of the SVM. The state can be Running, Stopped, Starting, Stopping, Not mapped, Initializing, or Deleting.

Volume count

Displays the number of volumes contained by the SVM.

Maximum Allowed Volumes

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

Root Volume

Displays the name of the root volume of the SVM.

Allowed protocols

Displays the type of protocols that can be configured on the SVM.

DNS Domain

Displays the DNS domain name.

NIS Domain

Displays the Network Information Service (NIS) domain name. This column is blank when the Network Information Service (NIS) server is disabled or is not configured.

LDAP Enabled

Displays if the LDAP protocol is enabled or not.

Name Service Switch

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

Related references

Storage Virtual Machine details page on page 185

Volume Data Protection Configuration report

The Volume Data Protection Configuration report enables you to view the unprotected volumes and Storage Virtual Machines (SVMs) that are used in a node or a cluster. This information enables you to understand the data protection risks for your system, and to view the details of the protected volumes and unprotected volumes in your system.

The Volume Data Protection Configuration report is displayed in two formats:

- Volume Data Protection Configuration report pie chart
- Volume Data Protection Configuration report tabular view

Volume Data Protection Configuration report pie chart

Displays the details of the protected volumes and unprotected volumes in your system.

Volume Data Protection Configuration report tabular view

Cluster

Displays the cluster name.

Storage Virtual Machine (SVM)

Displays the name of the Storage Virtual Machine (SVM) that contains the volume.

Volume

Displays the volume name.

Total Data Capacity (GB)

Displays the total data capacity (used plus available) in GB.

Used Data Capacity (GB)

Displays the used data capacity (in GB).

Used Data %

Displays the used data capacity as a percentage.

Available Data Capacity (GB)

Displays the available data capacity (in GB).

Available Data %

Displays the available data capacity as a percentage.

Snapshot Reserve Used Capacity (GB)

Displays the amount of space that is used by Snapshot copies from Snapshot reserve (in GB).

Snapshot Reserve Used %

Displays the amount of space that is used by Snapshot copies from Snapshot reserve as a percentage.

Snapshot Reserve Available Capacity (GB)

Displays the amount of space that is available for Snapshot copies (in GB).

Snapshot Reserve Available %

Displays the amount of space that is available for Snapshot copies as a percentage.

Snapshot Reserve Total Capacity (GB)

Displays the total snapshot reserve capacity of the aggregate (in GB).

Days To Full

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

Space Full Threshold %

Displays the percentage at which an aggregate is full.

Space Nearly Full Threshold %

Displays the percentage at which an aggregate is nearly full.

Daily Growth Rate %

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

Total Number Of Inodes

Displays the total number of inodes in the volume.

Inode Utilization

Specifies the inode space that is used in the volume.

Quota Committed Capacity

Displays the space that is reserved in the volumes.

Quota Overcommitted Capacity (GB)

Displays the amount of space that can be used (in GB) before the system generates the Volume Quota Overcommitted event.

Snapshot Autodelete

Displays whether automatic deletion of Snapshot copies is enabled or disabled.

Deduplication

Displays whether deduplication is enabled or disabled for the volume.

Deduplication Space Savings (GB)

Displays the amount of space that is saved in a volume by using deduplication (in GB).

Compression

Displays whether compression is enabled or disabled for the volume.

Compression Space Savings (GB)

Displays the amount of space that is saved in a volume by using compression (in GB).

Thin Provisioned

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

Autogrow

Displays whether the FlexVol volume automatically grows in size when it is out of space.

Space Guarantee

Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate.

State

Displays the state of the volume that is being exported.

Volume Relationships Inventory report

The Volume Relationships Inventory report enables you to analyze the storage inventory details in a cluster, understand the degree of protection that is required for volumes, and filter the volume details based on source of failure, pattern, and schedules.

The Volume Relationships Inventory report is displayed in two formats:

- Volume Relationships Inventory report pie chart
- Volume Relationships Inventory report tabular view

Volume Relationships Inventory pie chart

Displays the configuration details of the volume relationships that are present in your storage system.

Volume Relationships Inventory tabular view

Relationship Health

Displays the relationship heath of the cluster.

Relationship State
Displays the the mirror state of the SnapMirror relationship.

Relationship Status
Displays the status of the SnapMirror relationship.

Lag Status
Displays the lag status of the volume.

Source Cluster
Displays the name of the source cluster for the SnapMirror relationship.

Source SVM
Displays the name of the source Storage Virtual Machine (SVM) for the SnapMirror relationship.

Source Volume
Displays the name of the source volume for the SnapMirror relationship.

Destination Cluster
Displays the name of the destination cluster for the SnapMirror relationship.

Destination SVM
Displays the name of the destination Storage Virtual Machine (SVM) for the SnapMirror relationship.

Destination Volume
Displays the name of the destination volume for the SnapMirror relationship.

Volume Transfer Status report

The Volume Transfer Status report enables you to analyze the volume transfer trends over a period of time. You can configure the report to view the volume transfer status for a specific time interval. The report also displays whether the volume transfer was a success or a failure.

The Volume Transfer Status report is displayed in two formats:

- Volume Transfer Status report line chart view
- Volume Transfer Status report tabular view

Volume Transfer Status report line chart

The line chart displays the volume transfer details by plotting transfer count against date. You can also view whether a particular volume transfer has succeeded or failed.

Volume Transfer Status report tabular view

Source Cluster Name
Displays the source cluster name.

Source SVM
Displays the Storage Virtual Machine (SVM) name.

Source Volume Name
Displays the source volume name.

Destination Cluster Name
Displays the destination cluster name.

Destination SVM
Displays the destination SVM name.

Destination Volume Name

Displays the destination volume name.

Operation Result

Displays whether volume transfer was successful.

Start time

Displays the volume transfer start time.

End time

Displays the volume transfer end time.

Transfer duration (hrs)

Displays the time taken (in hours) to complete the volume transfer.

Transfer size (MB)

Displays the size (in MB) of the transferred volume.

Operation Type

Displays the type of volume transfer.

Volume Transfer Rate report

The Volume Transfer Rate report enables you to analyze the amount of data volume that is transferred on a day-to-day basis. The report also provides details about daily volume transfers and the time required to complete the transfer operation.

The Volume Transfer Rate report is displayed in two formats:

- Volume Transfer Rate report bar chart view
- Volume Transfer Rate report tabular view

Volume Transfer Rate report bar chart view

Displays the volume transfer rate details by plotting the total transfer size against the number of hours. You can also view the details of the amount of data that is transferred on a daily basis.

Volume Transfer Status report tabular view

Total Transfer Size (GB)

Displays the total size of the volume transfer in gigabytes.

Day

Displays the day on which the volume transfer was initiated.

End Time

Displays the volume transfer end time with date.

Schedule Report dialog box

You can schedule the reports to be generated on a recurring basis at a specified frequency from the Schedule Report dialog box. The report is sent by email to one or more users specified in the Schedule Report dialog box.

- *Properties* on page 456
- *Command buttons* on page 456

Properties

You can schedule a report by specifying properties such as the email address of the user, the format of the report, and the frequency at which the report is generated.

Using Existing Schedule

Schedule Name

Displays all the existing schedule names. You can select an existing schedule for your reports from here.

Create New Schedule

Schedule Name

Enables you to enter the schedule name while creating a new schedule.

Recipient Email Address

Specifies the email address of the user to whom you want to send the report. You can specify one or more entries, separated by commas. This is a mandatory field.

Report Format

Specifies the format in which you want to schedule the report. The *PDF* option is selected by default.

Frequency

Specifies the frequency at which you want to schedule the report. The *Hourly* option is selected by default.

Command buttons

The command buttons enable you to perform the following tasks:

Schedule

Schedules the report with the saved or updated template and closes the Schedule Report dialog box.

Cancel

Closes the Schedule Report dialog box while displaying a message to save the schedule report template.

Related concepts

What report scheduling is on page 416

Related tasks

Scheduling reports on page 411

Share Report dialog box

You can share a report with one or more users through email. After you customize a report, you must save the changes before you share the report to ensure that the changes are displayed.

- *Properties* on page 456
- *Command buttons* on page 457

Properties

You can share a report by specifying properties such as the email address of the user, subject of the email, and the format of the report.

Recipient Email Address

Specifies the email address of the user with whom you want to share the report. You can specify one or more entries, separated by commas. This is a mandatory field.

Subject

Specifies the subject of the email. By default, the name of the report is displayed.

Report Format

Specifies the format in which you want to share the report. The *PDF* option is selected by default. If the XHTML format is selected, open the report that is sent by email by using a supported web browser.

Command buttons

The command buttons enable you to perform the following tasks:

Share

Shares the report with the saved configuration and closes the Share Report dialog box.

Cancel

Closes the Share Report dialog box while displaying a message to save the report configuration.

Related concepts

What report sharing is on page 417

Related tasks

Sharing reports on page 411

Manage Report Schedules dialog box

You can view, modify, or delete existing report schedules and add new schedules for your reports from the Manage Report Schedules dialog box.

- *Properties* on page 456
- *Command buttons* on page 456

Properties

You can select an existing schedule or create a new schedule for your reports. You can view, modify, or delete your report schedules.

Left pane

Schedule Name

Displays the existing schedules. By clicking on any schedule you can view the schedule details in the right pane. For the first login, there are no existing schedules.

Add Schedule

Displays the new schedule form in the right pane. You can now add a new schedule.

Right pane

Schedule Name

Displays the schedule name.

Recipient Email Address

Displays the email address of the user to whom the report must be sent. You can enter more than one email addresses separated by commas.

Report Format

Displays the format in which the report must be presented. The PDF option is selected as the default report format. If the XHTML format is selected, open the report that is sent by email by using a supported web browser.

Frequency

Displays the frequency at which the report is scheduled.

Report Category

Displays the report category groups. Selecting a report category from the list, displays the reports that belong to that report category in the Available Reports column.

Available Reports

Displays only the reports that belong to the report category selected.

Selected Reports

Displays the selected reports to which you choose to apply the schedule. You can select the required reports from the Available Reports column. At least one report must be selected

Command buttons

The command buttons enable you to perform the following tasks:

Add Schedule

Enables you to add a new schedule.

Delete Schedule

Enables you to delete the schedule being currently viewed. When you create a new schedule, this button is not available.

Save

Saves the schedule being viewed, modified, or added.

Save and Close

Saves the schedule being viewed, modified, or added and closes the Manage Report Schedules dialog box.

Cancel

Closes the Manage Report Schedules dialog box while displaying a message to save the schedule.

Related concepts

What report scheduling is on page 416

Related tasks

Managing report schedules on page 412

Save Customized Report As dialog box

You can use the Save Customized Report As dialog box to save a report after customizing it.

- *Properties* on page 456
- *Command buttons* on page 457

Properties

You can customize and save a report by specifying properties such as the name and description.

Report Name

Displays the name of the report. The original report name is displayed by default. You can modify the report name as per the customization. Report name cannot exceed 255 characters.

Description

Specifies the description of the customization made on the report. Description cannot exceed 150 characters.

Command buttons

The command buttons enable you to perform the following tasks:

Save

Saves the customized report .

Cancel

Cancels the recent changes and closes the Save Customized Report As dialog box.

Related concepts

What reports do on page 414

Related tasks

Customizing a report on page 413

Save Custom Report dialog box

You can use the Save Custom Report dialog box to save a custom report after making additional changes to the custom report.

- *Properties* on page 456
- *Command buttons* on page 457

Properties

You can save a custom report by specifying properties such as the description.

Report Name

Displays the name of the custom report. This field cannot be edited.

Description

Specifies the description of the customization made on the custom report. Description cannot exceed 150 characters .

Command buttons

The command buttons enable you to perform the following tasks:

Save

Saves the custom report .

Cancel

Cancels the recent changes and closes the Save Custom Report dialog box.

Related concepts

What reports do on page 414

Related tasks

Editing a customized report on page 413

Import Report dialog box

You can use the Import Report dialog box to import reports from `.rptdesign` files.

Properties

You can import a report by specifying the report file name, report name, and report description.

Select Report File

Enables you to select the `.rptdesign` file that you want to import.

Note: In Google Chrome, the `fakepath` of the `.rptdesign` file is displayed. In Mozilla Firefox, only the `.rptdesign` file name is displayed. In Internet Explorer, the complete path of the `.rptdesign` file is displayed.

Name

Displays the name of the report. This field is empty by default. You can enter a name for the imported report.

Description

Specifies the description of the imported report. The description cannot exceed 150 characters.

Select database user with report schema role

Select, or create, a database user if you are importing reports from the Storage Automation Store.

Command buttons

The command buttons enable you to perform the following tasks:

Import

Validates the selected `.rptdesign` file, and imports the report.

Cancel

Cancels the import operation, and closes the Import Report dialog box.

Related concepts

What report importing is on page 417

Related tasks

Importing reports on page 414

Troubleshooting

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

Sending on-demand AutoSupport messages

You can configure OnCommand Unified Manager to on-demand messages to technical support for assistance with troubleshooting issues. The AutoSupport message contains diagnostic system information and detailed data about the Unified Manager server.

Before you begin

You must be logged in as the maintenance user.

Steps

1. Click ▦▾ > **Health**.
2. Click **Administration** > **Setup Options**.
3. In the **Setup Options** dialog box, click **Management Server** > **AutoSupport**.
4. Click **View AutoSupport Description**.

   The AutoSupport description is displayed, along with the product serial number, which is the number that technical support uses to find the AutoSupport messages that are sent from Unified Manager.
5. Perform one or both of the following actions:

If you want to send the AutoSupport message to...	Do this...
Technical support	Select the **Send to Technical Support** check box.
A specific email recipient	Select the **Send to Email Recipient** check box, and enter the email address of the recipient.

6. Click **Generate and Send AutoSupport**.

Related tasks

Adding a user on page 379

Enabling periodic AutoSupport on page 39

Unknown authentication error

Issue

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

1. From the Unified Manager Health page, click **Administration > Setup Options**.
2. In the Setup Options dialog box, click **Management Server > Authentication**.
3. Based on the authentication service that you have selected, enter the appropriate information for Administrator Name or Bind Distinguished Name.
4. Click **Test** to test the authentication with the details that you specified.
5. Click **Save and Close**.

User not found

Issue

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

1. From the Unified Manager Health page, click **Administration > Setup Options**.
2. In the Setup Options dialog box, click **Management Server > Authentication**.
3. Enter the appropriate information for base distinguished name.
4. Click **Save and Close**.

Icons are misaligned in Internet Explorer

Issue

Icons and text are misaligned when you use Internet Explorer.

Cause

This problem can occur if you are using Internet Explorer in Compatibility View, which is not a supported browser setting.

Corrective action

1. Press F12 to open Internet Explorer Developer Tools.
2. Select **Browser Mode** from the toolbar to display the browser version used to open the application.
3. Select **Document Mode** from the toolbar and select the Standards mode of the browser version used to open the application.

   For example, if you are using Internet Explorer 9 to open the application, select **Browser Mode > Internet Explorer 9**, and then select **Document Mode > Internet Explorer 9 Standards**.

LDAP server slow to respond

Issue

The LDAP server takes a long time to respond to queries.

Cause

Configuring the connection to support nested groups causes the LDAP server to slow down.

Corrective action

If you use Active Directory, you can speed up authentication by disabling support for nested groups in Unified Manager. If you choose to disable nested groups, you must ensure that users are direct members of the groups that are added to Unified Manager.

To disable nested group support, follow these steps:

1. From the top-level Health page, click **Administration > Setup Options**.
2. In the Setup Options dialog box, click **Management Server > Authentication**.
3. From the **Authentication Service** drop-down list, select **Others**.
4. In the **Member** box, type **`member`**.
5. Click **Save and Close**.

Issue with adding LDAP using Other authentication services

Issue

When you select Other in the Authentication section of the Setup Options dialog box, the user and groupObjectClass 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:

1. Click ▦▾ > **Health**.
2. Click **Administration > Setup Options**.
3. In the Setup Options dialog box, click **Management Server > Authentication**.
4. In the **Authentication Service** drop-down menu, select **Active Directory**, and then select **Others**.
5. Complete the text fields.

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

1. Click ▦▾ > **Health**

2. Click **Administration > Setup Options**.
3. In the Setup Options dialog box, click **Management Server > Authentication**.
4. In the **Authentication Service** drop-down menu, select **OpenLDAP**, and then select **Others**.
5. 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.

Troubleshooting access to CIFS shares

You might not be able to access CIFS shares if the storage objects serving these shares are unavailable. You should review availability events such as Volume Offline, Junction Path Offline, or SVM CIFS Server Down that are generated when these objects are unavailable.

Before you begin

You must have the role of Storage Administrator to perform this task.

About this task

If you have configured an appropriate alert, you will be notified about the availability event through an alert email.

Steps

1. Log in to Unified Manager UI.
2. From the **Dashboard**, click the appropriate offline event.

   Example

   For example, if you receive a Volume Offline event, click the ***`Volume_name`*** **Volume Offline** event in the Availability panel in the Unresolved Incidents and Risks area.
3. In the **Event details** page, click ***`Volume_name`*** in the **Source** field.
4. In the **Volume details** page, click the number corresponding to CIFS Shares in the **Related Devices** pane.
5. In the **Storage Virtual Machine details** page, click the **CIFS Shares** tab.

   You can view the number of CIFS shares that are affected.

After you finish

You must resolve the failures by using either OnCommand System Manager or the Data ONTAP CLI.

Identifying inaccessible volumes and qtrees

You might not be able to access a volume or a qtree or write data to the volume or qtree. In such cases, the storage administrator must identify the quota for the volume or qtree that is not accessible and take appropriate action to resolve the issue.

Before you begin

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

Steps

1. Log in to Unified Manager.
2. In the search criteria, select **User or Group Quotas**, enter the user name, and press Enter.

   #### Example

   If a user with user name user1 cannot access a volume, select **User or Groups Quotas** in the search criteria, and then type **`user1`** or type the any three characters of the user name in the search bar.
3. Select the appropriate user from the drop-down list.

   The User and Group Quotas tab in the Storage Virtual Machine details page is displayed.

After you finish

You must resolve the issue by using either OnCommand System Manager or the Data ONTAP CLI.

Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring

Issue

If performance monitoring is enabled on the Unified Manager server through connections with one or more Performance Manager servers, and if an HTTPS certificate is installed or regenerated on that Unified Manager server and that server is rebooted, posting of additional performance incidents encountered by Performance Manager servers to the Unified Manager Web UI is stopped.

Cause

The HTTPS certificate installation or regeneration on the Unified Manager server has invalidated the credentials that allowed the Performance Manager servers to post performance incidents to the Unified Manager server.

Corrective actions

You need to delete and then reconfigure the connections between the Unified Manager server and each Performance Manager server that was posting performance incidents to it before the HTTPS certificate installation or regeneration.

Before starting corrective actions, be prepared to respecify the current connection information:

- Unified Manager server name or IP address
- Unified Manager server port (always 443)

- Event Publisher user name (the name of the local Unified Manager server user assigned Event Publisher role privileges)
- Event Publisher password (the password of the local Unified Manager server user assigned Event Publisher role privileges)

Log in to the maintenance console of each Performance Manager server that was posting performance incidents to the Unified Manager server, and then complete the following actions:

1. Type the number of the menu option labeled "Unified Manager Connection" to display the Unified Manager Connection Menu, and then type the number of the menu option labeled "Delete Unified Manager Server Connection."
2. When prompted to confirm that you want to continue the deletion, type **y**, and then press any key to continue.
3. Type the number of the menu option labeled "Add / Modify Unified Manager Server Connection."
4. When prompted, supply the requested Unified Manager server name or IP address and Unified Manager server port information.
5. When prompted, accept the Unified Manager server trust certificate to support the connection. Do not change the default port value (443).
6. When prompted, supply the requested Event Publisher user name and Event Publisher password, and then confirm that the settings are correct.
7. To exit the maintenance console, press any key to continue, and then type **x**.

Certain special characters do not work with reporting search

Issue

Using the special characters % and _ while searching within a report causes the operation to fail.

Corrective action

If you search for a string that contains % or _, you should use a double backslash before the specified character.

For example, to find a string containing S_10, you should enter S\\_10.

Glossary

A

Access Control List (ACL)

A set of data associated with a file, directory, or other resource or share that defines user or group access rights to that resource or share.

admin SVM

Formerly known as admin Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that has overall administrative access to all objects in the cluster, including all objects owned by other SVMs, but does not provide data access to clients or hosts.

aggregate

A set of multiple RAID (Redundant Array of Inexpensive Disks or Redundant Array of Independent Disks) groups that can be managed as a single unit for protection and provisioning purposes.

aggregate committed capacity

The data storage space in an aggregate that is committed to provide for its underlying volumes. Calculated by the total capacity provisioned for volumes.

aggregate total capacity

The data storage space within an aggregate that can be used by volumes or aggregate-level Snapshot copies. Calculated by the total data capacity of the aggregate plus the aggregate-level Snapshot reserve space.

alert

- In OnCommand Insight (formerly SANscreen suite), an alarm indicating that a device or path state has changed in a way that violates a policy or exceeds a threshold.
- In Unified Manager, a user-configured notification that is sent whenever a specific event or an event of a specific severity type occurs, not necessarily related to a specific user. Alerts are used to monitor and manage datasets and resources. See also *event* and *severity type*.

AutoSupport

An integrated technology that triggers email messages from the customer site to technical support or another specified email recipient when there are any failures in Unified Manager services. These messages contain information such as feature usage metrics, configuration and user settings, system health, and so on.

available capacity

The amount of usable space available in a storage system. Calculated by the used capacity minus the unused reserve capacity.

B

backup relationship

A persistent association between a primary directory and a secondary volume for disk-based data backup and restore using the Data ONTAP SnapVault feature.

baseline transfer

An entire transfer of data as compared to an incremental transfer of data.

C

CIFS

See *Common Internet File System (CIFS)*.

CIFS share

- In Data ONTAP, a directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known simply as a *share*.
- In OnCommand Insight (formerly SANscreen suite), a service exposed from a NAS device to provide file-based storage through the CIFS protocol. CIFS is mostly used for Microsoft Windows clients, but many other operating systems can access CIFS shares as well.

client application

An application that calls Unified Manager APIs to enable its operator to configure, monitor, and initiate data management operations to be executed on the Unified Manager server.

cluster

- In clustered Data ONTAP 8.x, a group of connected nodes (storage systems) that share a namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.
- In the Data ONTAP 7.1 release family and earlier releases, a pair of storage systems (sometimes called *nodes*) configured to serve data for each other if one of the two systems stops functioning.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.
- For some storage array vendors, *cluster* refers to the hardware component on which host adapters and ports are located. Some storage array vendors refer to this component as a *controller*.

cluster committed capacity

The data storage space in a cluster that is committed to provide for its underlying aggregates. Calculated by the sum of the total capacity of all the aggregates in the cluster.

cluster failover (CFO)

In Data ONTAP 7.1.x and earlier, the method of ensuring data availability by transferring the data service of a failed node to another node in an active/active configuration. Transfer of data service is often transparent to users and applications. In Data ONTAP 7.2 and later, and in Data ONTAP operating in 7-Mode, the failover method is called *controller failover*. In clustered Data ONTAP, the failover method is called *storage failover*.

cluster interconnect

The cables and adapters with which two nodes (storage systems) in an HA pair are connected, and over which heartbeat and WAFL log information are transmitted when both nodes are running.

cluster total capacity

The data storage space in a cluster that can be used by aggregates or volumes. Calculated by the sum of the capacity of all the data disks excluding disk right-sizing and reservation plus sum of the capacity of all spare disks excluding right-sizing.

cluster Vserver

Former name for a data SVM; see data SVM.

container object

An object, such as an aggregate or a Storage Virtual Machine (SVM, formerly known as Vserver), in which data objects reside.

counter

The statistical measurement of activity on a storage system or storage subsystem that is provided by Data ONTAP. Each type of storage system or subsystem has a set of counters.

D

Data ONTAP

The operating system software running on NetApp storage devices.

datastore

A storage location for virtual machines, such as a VMFS volume, a directory on a NAS server, or a local file system path. A datastore is platform-independent and host-independent; therefore, it does not change when a virtual machine it contains moves to another host.

data object

A container of data, such as a file, directory, volume, or LUN, that can be discovered, monitored, protected, created, or restored by the Unified Manager server.

data SVM

Formerly known as data Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data SVMs within a cluster.

dedicated cache SSDs

The disks that are assigned to an aggregate and used for cache. These disks are not shared among aggregates.

dedicated data disks

The disks that are assigned to an aggregate and used for storing data. These disks are not shared among aggregates.

deduplication

The consolidation of blocks of duplicate data into single blocks to store more information using less storage space.

deduplication return

The capacity savings resulting from deduplication. Calculated by the volume capacity before deduplication - the volume capacity after deduplication.

destination

The storage to which source data is backed up, mirrored, or migrated.

destination data object

A data object that contains the backed up or mirrored replicated data.

dedupe

See *deduplication*.

DHCP

See *Dynamic Host Configuration Protocol (DHCP)*.

Dynamic Host Configuration Protocol (DHCP)

The protocol for automating the assignment of network addresses.

E

ESX server

A VMware term describing a server that abstracts server processor, memory, storage, and networking resources into multiple virtual machines.

event

An indication of a predefined condition occurring or when an object crosses a threshold. All events are assigned a severity type and are automatically logged in the Events window. See also *alert* and *severity type*.

F

failover

The process by which an alternate storage system takes over and emulates a primary system if the primary system becomes unusable.

Fibre Channel (FC)

A high-speed data transmission protocol, which is a licensed service on the storage system that enables you to export LUNs to hosts using the SCSI protocol over an FC fabric.

FlexVol volume

In clustered Data ONTAP, a logical entity contained in a Storage Virtual Machine (SVM, formerly known as Vserver)—referred to as SVM with FlexVol volumes. FlexVol volumes typically hold user data, although they also serve as node or SVM root volumes and metadata containers. A FlexVol volume obtains its storage from a single aggregate.

FQDN

See *Fully Qualified Domain Name (FQDN)*.

Fully Qualified Domain Name (FQDN)

The complete name of a specific computer on the Internet, consisting of the computer's host name and its domain name.

fractional reserve

An option that determines how much space in a volume is reserved for Snapshot overwrite data for LUNs and space-reserved files, to be used after all other space in the volume is used.

G

giveback

The return of identity from an emulated storage system to the failed system, resulting in a return to normal operation. The reverse of *takeover*.

global namespace

See *namespace*.

growth rate

The measurement of how fast the storage is filling. The growth rate is determined by dividing the daily growth rate by the total amount of space in the storage system.

H

HA (high availability)

- In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an *HA pair*, that are configured to serve data for each other if one of the two nodes stops functioning.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

HA pair

- In Data ONTAP 8.x, a pair of nodes whose controllers are configured to serve data for each other if one of the two nodes stops functioning.
  Depending on the system model, both controllers can be in a single chassis, or one controller can be in one chassis and the other controller can be in a separate chassis.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

host

A computer system that accesses data on a storage system.

host bus adapter (HBA)

An interface card that plugs into a SAN device. SAN devices use the ports on their respective HBAs to connect to each other in a SAN. Each SAN device might contain one or more HBAs, and an HBA might contain more than one port. Each port can be used to establish a connection to a SAN.

I

igroup

initiator group. A collection of unique iSCSI node names of initiators (hosts) in an IP network that are given access to *front-end LUNs* when they are mapped to those LUNs. (Array LUNs on a storage array that provide storage for Data ONTAP systems can be considered *back-end LUNs*.)

incident

An issue that has already impacted the availability or capacity of storage objects.

incremental transfer

A subsequent backup after a baseline transfer has occurred of a primary directory in which only the new and changed data since the last backup (baseline or incremental) is transferred. The transfer time of incremental transfers can be significantly less than the baseline transfer.

Infinite Volume

In clustered Data ONTAP, a logical entity contained in a Storage Virtual Machine (SVM, formerly known as Vserver)—referred to as SVM with Infinite Volume—that holds user data. An Infinite Volume obtains its storage from multiple aggregates.

initiator

The system component that originates an I/O command over an I/O bus or network. The target is the component that receives this command.

inode

A data structure containing information about files on a storage system and in a UNIX file system.

inter-switch link (ISL)

A connection between two switches using the E-port.

iSCSI

Internet Small Computer Systems Interface (iSCSI) protocol. A licensed service on the storage system that enables you to export LUNs to hosts using the SCSI protocol over TCP/IP.

iSCSI router

A storage router implementing the Internet Small Computer Systems Interface (iSCSI) protocol (SCSI over IP) to extend access of a Fibre Channel fabric and attached storage devices to IP servers.

J

JBOD

Just a Bunch Of Disks. An array of disks without any redundancy; that is, without RAID configuration.

job

A long-running operation, for example, scheduled local backup of a dataset, a mirror transfer, and password updates.

L

level-0 backup

An initial backup (also known as a *baseline transfer*) of a primary directory to a secondary volume in which the entire contents of the primary directory are transferred.

LIF (logical interface)

A logical network interface, representing a network access point to a node. LIFs currently correspond to IP addresses, but could be implemented by any interconnect. A LIF is generally bound to a physical network port; that is, an Ethernet port. LIFs can fail over to other physical ports (potentially on other nodes) based on policies interpreted by the LIF manager.

Lightweight Directory Access Protocol (LDAP)

A client-server protocol for accessing a directory service.

local backup

Local backup protection (also referred to as *Snapshot protection*) is the periodic capture of the active data on a NetApp storage system in backup images and the storage of those images on that same system. If active data on the local system is accidentally deleted or corrupted, it can quickly be restored with the most recent image stored locally from the last local backup job. Local backup operations are typically employed on the primary storage systems, where data is being actively updated and where, in event of accidental data loss, data restoration from the last hour or two might be required. Local backup protection is based on NetApp Snapshot technology.

local backup copy

A copy of data, usually on a primary node, created using Snapshot technology and that resides on the primary dataset node.

Logical Unit Number (LUN)

A SCSI identifier of a logical unit of storage on a target. LUNs are often referred to as *virtual disks*, and vice versa. See also *virtual disk*.

logical object

The entity that represents a storage container, such as a volume, qtree, LUN, or dataset.

lower threshold

The value set to generate an event when a counter falls and remains below that value for longer than the specified interval.

M

maintenance user

The user who has access rights to deploy and configure an OnCommand Unified Manager virtual appliance.

Management Information Base (MIB)

ASCII files that describe the information that the SNMP agent sends to network management stations.

member

Any data object that subscribes to or is created by a storage service.

mirror (v)

The process of creating an exact duplicate of all volume data from a NetApp source storage system to a destination storage system. If data in the source storage system is lost or made unavailable, then that replicated data can quickly be made available from the destination mirror site. Mirror operations are employed from primary to secondary storage and from secondary to tertiary storage. Mirror protection is based on NetApp Volume SnapMirror technology.

mirror copy (n)

The exact duplicate of all volume data (both active and protected) from a NetApp source storage system to a destination storage system, created using NetApp Volume SnapMirror technology.

move (v)

To physically move data and any needed associated configuration of an object from one aggregate to another within a cluster, including within a single node.

N

namespace

In network-attached storage (NAS) cluster environments, an abstraction layer for data location that provides a single access point for all data in the system. It enables users to access data without specifying the physical location of the data, and enables administrators to manage distributed data storage as a single file system. Sometimes referred to as *global namespace*.

NDMP

Network Data Management Protocol. A protocol that allows storage systems to communicate with backup applications and provides capabilities for controlling the robotics of multiple tape backup devices.

Network File System (NFS) export

A service exposed from a NAS device to provide file-based storage through the NFS protocol. NFS is mostly used for UNIX-like operating systems, but other operating systems can access NFS exports as well.

node

- In Data ONTAP, one of the systems in a cluster or an HA pair.
  To distinguish between the two nodes in an HA pair, one node is sometimes called the *local node* and the other node is sometimes called the *partner node* or *remote node*.
- In Protection Manager and Provisioning Manager, the set of storage containers (storage systems, aggregates, volumes, or qtrees) that are assigned to a dataset and designated either primary data (primary node), secondary data (secondary node), or tertiary data (tertiary node).
  A *dataset node* refers to any of the nodes configured for a dataset.
  A *backup node* refers to either a secondary or tertiary node that is the destination of a backup or mirror operation.
  A *disaster recovery node* refers to the dataset node that is the destination of a failover operation.

nondisruptive

The ability of a system to continue serving data to clients during a system process or activity, such as a LUN restore operation or an online migration.

NVRAM mirror

A synchronously updated copy of the contents of the storage system NVRAM (nonvolatile random access memory) contents kept on the partner storage system.

O

offline

A database state indicating that the database is not available to users (for example, the database is in the following states: shutdown, started, or mounted).

online

A database state indicating that the database is available to users (for example, open).

OnCommand administrator

An RBAC role that enables a person to configure settings for items unrelated to storage management, such as user roles, security certificates, database access, LDAP, SMTP, networking, and AutoSupport.

operator

An RBAC role that enables a person to view data and to view, assign, and resolve events in OnCommand Unified Manager.

P

parity disk

The disk on which parity information is stored for a RAID4 disk drive array. In RAID groups using RAID-DP protection, two parity disks store the parity and double-parity information. Used to reconstruct data in failed disk blocks or on a failed disk.

partner node

From the point of view of the local node (storage system), the other node in an HA configuration.

plex

A collection of one or more RAID groups that together provide the storage for one or more Data ONTAP volumes.

port

A physical connection point on computers, switches, storage arrays, and so on, which is used to connect to other devices on a network. Ports on a Fibre Channel network are identified by their World Wide Port Name (WWPN) IDs. TCP/IP ports are used as virtual addresses assigned to each IP address.

protection artifact

An object, such as a destination data object, or a protection relationship that the Unified Manager server creates to support protection jobs when a data object is subscribed to a storage service.

protection policies

The entities that enable you to set the automation controls for scheduling, monitoring, and alerts on any set of data in terms of normal backup, offsite backup, disaster and recovery backup, and regulatory copies.

protection relationship

The SnapMirror or SnapVault relationship that exists between a source data object and a destination data object.

Q

qtree

A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.

R

RAID-DP

redundant array of independent disks, double-parity.

raw capacity

The total amount of addressable blocks on physical disk drives. Calculated by multiplying the number of disk drives by the labeled capacity of those disk drives. For Data ONTAP systems using array LUNs, it refers to the size of the array LUNs.

RBAC

Role-based Access Control. A system whereby access to resources is decided based on the role of a user. RBAC controls who has access to various operations on which resources. Access to resources is first assigned to roles and roles are then assigned to users. Conforms to NIST RBAC standard.

recovery

The re-creation of a past operational state of an entire application or computing environment. Compare *restore*. A recovery operation can encompass a restore. *Recovery* and *restore* are often used synonymously. Recovery and restore can also be *in-place*, meaning that copies are mounted and used locally instead of being copied elsewhere.

remote backup

A copy of data on another set of physical disks or medium. Also referred to as *secondary storage*. See also *local backup*.

remote storage

The storage that is accessible to the local node, but is at the location of the remote node.

replication

The process of duplicating data from one highly available site to another. The replication process can be synchronous or asynchronous. Duplicates are known as *clones*, *point-in-time copies*, or *Snapshot copies*, depending on the type of copy being made.

restore

The copying of an object, such as a file or an attribute, or an entire application or virtual machine, back to its original source. Compare *recovery*. A restore can be part of a recovery operation. *Restore* and *recovery* are often used synonymously. Restore and recovery can also be *in-place*, meaning that copies are mounted and used locally instead of being copied elsewhere.

retention period

The user-specified minimum length of time that a local backup Snapshot copy must be retained.

root member

A data object that subscribes to a storage service.

S

SAN

storage area network. A network linking servers or workstations to devices, typically over Fibre Channel. The SAN model places storage on its own dedicated network, removing data storage from both the server-to-disk SCSI bus and the main user network. The SAN includes one or more hosts that provide a point of interface with LAN users, as well as one or more fabric switches (in the case of large SANs) and SAN hubs to accommodate a large number of storage devices.

severity type

The specific severity type of an event when an event occurs. Each event is associated with a severity type to help you determine priorities for taking corrective action.

SFO

See *storage failover (SFO)*.

share

A directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known as a *CIFS share*.

shared cache SSDs

The disks that are assigned to an aggregate and used for cache. These disks are shared among aggregates.

shared data disks

The disks that are assigned to an aggregate and used for storing data. These disks are shared among aggregates.

Snapshot available capacity

The storage space that is available in the Snapshot reserve for Snapshot copies. Calculated by subtracting the total Snapshot used capacity from the Snapshot reserve capacity.

Snapshot copy

An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copies enable users to restore files and to back up the storage system to tape while the storage system is in use.

Snapshot overflow

The storage space that is consumed by Snapshot copies from the total data capacity of a volume or an aggregate. Calculated by subtracting the Snapshot reserve capacity from the Snapshot used capacity.

Snapshot reserve capacity

The storage space that is set aside by the volume or the aggregate for its Snapshot copies. Data cannot be written to this space.

Snapshot return

The capacity savings of Snapshot copies when compared to full volume copies. Calculated by the volume capacity - the Snapshot capacity.

Snapshot unused capacity

The capacity remaining after some capacity that was reserved for Snapshot copies is used.

Snapshot used capacity

The capacity that has been used for Snapshot copies.

spare disk

A physical disk that is part of a storage device that is the same technology type (FC, SATA), size, and speed as a standard disk. A spare disk is used in case the standard disk malfunctions.

standard HA configuration

A configuration set up so that one node automatically takes over for its partner when the partner node becomes impaired.

storable capacity

The disk capacity that is available for use after right-sizing.

storage administrator

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

storage controller

The component of a storage system that runs the operating system and controls its disk subsystem. Storage controllers are also sometimes called *controllers*, *storage appliances*, *appliances*, *storage engines*, *heads*, *CPU modules*, or *controller modules*.

storage efficiency

The ratio of usable capacity to effective used capacity, accounting for efficiency returns. Calculated by the effective used capacity / the usable capacity.

storage failover (SFO)

The method of ensuring data availability by transferring the data service of a failed node to another node in the cluster. Transfer of data service is often transparent to users and applications. Also referred to as *controller failover (CFO)*or *cluster failover (CFO)*.

storage utilization

The ratio of usable capacity to used capacity, without accounting for efficiency returns. Calculated by the used capacity / the usable capacity.

Storage Virtual Machine (SVM)

(Known as *Vserver* prior to clustered Data ONTAP 8.2.1. The term “Vserver” is still used in CLI displays and `vserver` command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs—*admin*, *node*, and *data*—but unless there is a specific need to identify the type of SVM, “SVM” usually refers to the data SVM.

SVM

(Storage Virtual Machine; known as *Vserver* prior to clustered Data ONTAP 8.2.1. The term “Vserver” is still used in CLI displays and `vserver` command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs—*admin*, *node*, and *data*—but unless there is a specific need to identify the type of SVM, “SVM” usually refers to the data SVM.

SVM guaranteed available capacity

The data storage space that is guaranteed by the SVM to its underlying volumes but is not used. Calculated as the sum of the available size of all the thick provisioned volumes.

SVM used capacity

The data storage space that is guaranteed by the SVM to its underlying volumes. Calculated as the sum of the used data capacity of all the volumes associated with the SVM.

SVM unguaranteed capacity

The data storage space that is not guaranteed by the SVM to its underlying volumes. Calculated as the sum of the available sizes of all the thin provisioned volumes.

SVM total capacity

The sum of the total data storage space of all the volumes in the SVM.

SVM unguaranteed capacity

The data storage space that is not guaranteed by the SVM to its underlying volumes. Calculated as the SVM total capacity minus the sum of the committed capacity of aggregates that are associated with the SVM.

switchback

The MetroCluster operation that restores service back to one of the MetroCluster sites.

switchover

The MetroCluster operation that transfers service from one of the MetroCluster sites.

- A *negotiated* switchover is planned in advance and cleanly shuts down components of the target MetroCluster site.
- A *forced* switchover immediately transfers service; the shut down of the target site might not be clean.

system reserve capacity

The capacity required for fixed system reserves, RAID parity, mirroring, and spare drives. Calculated by the fixed reserve + RAID reserve + spares.

T

takeover

The emulation of the failed node identity by the takeover node in an HA pair; the opposite of *giveback*.

takeover node

A node (storage system) that remains in operation after the other node stops working and that hosts a virtual node that manages access to the failed node disk shelves and network connections. The takeover node maintains its own identity and the virtual node maintains the failed node identity.

thin provisioning (TP)

A method of optimizing the efficiency with which the available space is used in storage. When many applications share access to the same storage array, thin provisioning enables administrators to maintain a single free space buffer pool to service the data requirements of all applications. This is done by allocating disk storage space in a flexible manner among multiple consumers, based on the minimum space required by each at any given time. This avoids poor utilization that occurs on traditional storage arrays where large pools of storage capacity are allocated to individual applications, but much remains unused.

trap

An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the storage system.

U

unassigned disk

A disk that is not assigned to any node or counted as a spare disk.

unused capacity

The free, usable storage space available for storing user data on the device, excluding capacity reserved for Snapshot copies or data.

unused reserve capacity

The capacity allocated but unused by Aggregate Snapshot Reserve, Volume Snapshot Reserve, Volume Fractional Reserve, and Vol/LUN/File Guaranteed Space. These reserves are adjustable by the user. Calculated by the aggregate snapshot unused reserve + volume snapshot unused reserve + volume fractional unused reserve + vol/LUN/file unused guaranteed space.

usable capacity

The capacity available to applications and users. Calculated by the raw capacity - the system reserve capacity.

used capacity

The capacity used by application or user data, including volume Snapshot copies and aggregate Snapshot copies. Calculated by the usable capacity - the free capacity.

V

vApp

See *virtual appliance*.

virtual appliance

A prebuilt software solution containing virtual machines and software applications that are integrated, managed, and updated as a package. Also called *vApp*.

virtual machine

A guest operating system and any application installed thereon, running on a computing device on which the software is installed, or suspended to disk or any other storage media accessible by the computing device.

VMware VirtualCenter (VC)

A management software suite used to create your VMware datastores and virtual machines and to configure the storage system volumes as the containers in which your active datastore and virtual machine images reside. It consists of the following components:

- VMware agents—software modules installed on an ESX server to carry out VC server requests
- VirtualCenter (VC) server—a server communicating with VMware agents on an ESX server
- Virtual Infrastructure (VI) client—a GUI client to manage the VC server
- VMware ESX server—an enterprise-level product that integrates server processes, storage functionality, and networking resources into multiple virtual systems.

  ESX server can also refer to a physical host running an ESX server hypervisor OS.

volume

- For Data ONTAP, a logical entity that holds user data that is accessible through one or more of the supported access protocols, including Network File System (NFS), Common Internet File System (CIFS), Fibre Channel (FC), and Internet SCSI (iSCSI). Data ONTAP treats an IBM volume as a disk.
- For IBM, the area on the storage array that is available for a Data ONTAP system or non Data ONTAP host to read data from or write data to. The documentation uses the term *array LUN* to describe this area.

volume committed capacity

The data storage space in a volume that is committed to provide storage space for its underlying qtrees based on their quota settings. Calculated as the sum of all qtrees disk hard limits. The sum does not include the qtrees for which the disk hard limit is not set.

volume total capacity

The data storage space in a volume that can be used by qtrees, LUNs, or other files and volume-level Snapshot copies. Calculated as the total data capacity of the volume plus the volume-level Snapshot reserve space.

Vserver

(Known as “Storage Virtual Machine (SVM)” in clustered Data ONTAP 8.2.1 and later.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of Vservers—*admin*, *node*, and *cluster* (“cluster Vserver” is called “data Vserver” in Data ONTAP 8.2)—but unless there is a specific need to identify the type of Vserver, “Vserver” usually refers to the cluster/data Vserver.

Copyright information

Trademark information

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index