

OnCommand® Unified Manager 7.3

Workflow Guide for Managing Cluster Health

November 2020 | 215-12453_2020-11_en-us
doccomments@netapp.com

NetApp®

Contents

Introduction to OnCommand Unified Manager health monitoring..... 5
Unified Manager health monitoring features....5
Unified Manager interfaces used to manage storage system health....6
OnCommand Unified Manager product documentation....7

Common Unified Manager health workflows and tasks..... 8
Configuring your environment after deployment....10
Changing the Unified Manager virtual appliance host name....11
Changing the Unified Manager host name in Red Hat Enterprise Linux....12
Adding clusters....14
Configuring Unified Manager to send alert notifications....15
Configuring database backup settings....27
Changing the local user password....28
Monitoring and troubleshooting data availability....28
Resolving a flash card offline condition....28
Scanning for and resolving storage failover interconnect link down conditions....30
Resolving volume offline issues....32
Resolving capacity issues....36
Performing suggested remedial actions for a full volume....37
Creating, monitoring, and troubleshooting protection relationships....38
Setting up protection relationships in Unified Manager....38
Performing a protection relationship failover and failback....45
Resolving a protection job failure....49
Resolving lag issues....52
Restoring data from Snapshot copies....54
Restoring data using the Health/Volume details page....54
Restoring data using the Health/Volumes inventory page....55
Managing scripts....55
How scripts work with alerts....56
Adding scripts....57
Deleting scripts....57
Testing script execution....58
Managing and monitoring groups....58
Understanding groups....59
Adding groups....63
Editing groups....63
Deleting groups....64
Adding group rules....64
Editing group rules....66
Deleting group rules....66
Adding group actions....67
Editing group actions....67
Configuring volume health thresholds for groups....68
Deleting group actions....69
Reordering group actions....69
Prioritizing storage object events using annotations....70

Understanding more about annotations....71
Adding annotations dynamically....73
Adding values to annotations....74
Deleting annotations....74
Viewing the annotation list and details....75
Deleting values from annotations....75
Creating annotation rules....75
Adding annotations manually to individual storage objects....77
Editing annotation rules....77
Configuring conditions for annotation rules....78
Deleting annotation rules....79
Reordering annotation rules....79
Configuring backup and restore operations....80
What a database backup is....80
Configuring database backup settings....80
What a database restore is....81
Virtual appliance backup and restore process overview....82
Restoring a database backup on a virtual machine....82
Restoring a database backup on Red Hat Enterprise Linux....83
Restoring a database backup on Windows....84
Managing SAML authentication settings....85
Identity provider requirements....85
Enabling SAML authentication....86
Changing the identity provider used for SAML authentication....87
Updating SAML authentication settings after Unified Manager security certificate change....88
Disabling SAML authentication....89
Disabling SAML authentication from the maintenance console....90
Managing storage objects using the Favorites option....91
Adding to, and removing storage objects from, the Favorites list....91
Cluster favorite card....92
Aggregate favorite card....92
Volume favorite card....93
Creating and importing reports into Unified Manager....94
Downloading and installing MySQL Connector/J....94
Creating a database user....95
Downloading the Eclipse Business Intelligence and Reporting Tools (BIRT)....95
Creating a project using BIRT....95
Creating a new report using BIRT....96
Creating a JDBC data source using BIRT....96
Creating a new MySQL data set using BIRT....97
Importing reports....97
Setting up and monitoring an SVM with Infinite Volume without storage classes....97
Editing the Infinite Volume threshold settings....99
Managing your Infinite Volume with storage classes and data policies....99
Editing the threshold settings of storage classes....100
Adding alerts....101
Creating rules....103
Exporting a data policy configuration....104
Sending a Unified Manager support bundle to technical support....105
Accessing the maintenance console....106
Generating a support bundle....106
Retrieving the support bundle using a Windows client....108
Retrieving the support bundle using a UNIX or Linux client....108
Sending a support bundle to technical support....109

Related tasks and reference information....109
Adding and reviewing notes about an event....110
Assigning events to specific users....110
Acknowledging and resolving events....110
Event details page....111
Description of event severity types....114
Description of event impact levels....115
Description of event impact areas....115
Health/Volume details page....116
Health/Storage Virtual Machine details page....128
Health/Cluster details page....142
Health/Aggregate details page....153
Protection/Job details page....161
Definitions of user roles....162
Definitions of user types....163
Unified Manager user roles and capabilities....163
Supported Unified Manager CLI commands....165

Using the maintenance console....170
What functionality the maintenance console provides....170
What the maintenance user does....170
Diagnostic user capabilities....171
Accessing the maintenance console....171
Accessing the maintenance console using the vSphere VM console....172
Maintenance console menus....172
Network Configuration menu....173
System Configuration menu....174
Support and Diagnostics menu....175
Additional menu options....175
Changing the maintenance user password on Windows....176
Changing the umadmin password on Red Hat Enterprise Linux....176
Adding network interfaces....177
Adding disk space to the Unified Manager database directory....177
Adding space to the data directory of the Red Hat Enterprise Linux host....178
Adding space to the data disk of the VMware virtual machine....179
Adding space to the logical drive of the Microsoft Windows server....179

Copyright and trademark....181
Copyright....181
Trademark....181

Introduction to OnCommand Unified Manager health monitoring

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

The key capabilities of Unified Manager include monitoring, alerting, managing availability and capacity of clusters, managing protection capabilities, 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.

Unified Manager health monitoring features

Unified Manager is built on a server infrastructure that delivers scalability, supportability, and enhanced monitoring and notification capabilities. Unified Manager supports monitoring of systems running ONTAP software.

Unified Manager includes the following features:

- Discovery, monitoring, and notifications for systems that are installed with ONTAP software:
  - Physical objects: 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, FabricPool aggregates, deduplication, and compression
  - Protection: SnapMirror relationships and SnapVault relationships
- 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 alerts, events, and threshold infrastructure
- LDAP, LDAPS, SAML authentication, and local user support
- RBAC (for a predefined set of roles)
- AutoSupport and support bundle
- 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:
  - FlexVol volumes
  - FlexGroup volumes
  - Infinite Volumes
- 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, space savings, and related events
- Integration with OnCommand Workflow Automation to execute workflows

  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*

Unified Manager interfaces used to manage storage system health

This guide contains information about the two user interfaces that OnCommand Unified Manager provides for troubleshooting data storage capacity, availability, and protection issues. The two UIs are the Unified Manager web UI and the maintenance console.

If you want to use the protection features in Unified Manager, you must also install and configure OnCommand Workflow Automation (WFA).

Unified Manager web UI

The Unified Manager web UI enables an administrator to monitor and troubleshoot cluster issues relating to data storage capacity, availability, and protection.

This guide describes some common workflows that an administrator can follow to troubleshoot storage capacity, data availability, or protection issues displayed in the Unified Manager web UI.

Maintenance console

The maintenance console enables an administrator to monitor, diagnose, and address operating system issues, version upgrade issues, user access issues, and network issues related to the Unified Manager server itself. If the Unified Manager web UI is unavailable, the maintenance console is the only form of access to Unified Manager.

This guide provides directions for accessing the maintenance console and using it to resolve issues related to the functioning of the Unified Manager server.

OnCommand Unified Manager product documentation

OnCommand Unified Manager is accompanied by a set of guides that describe how to install and use the product. Online help is also provided in the user interface.

OnCommand Unified Manager Installation and Setup Guide

Provides installation, upgrade, and setup instructions for Unified Manager on the VMware, Red Hat, and Windows platforms.

OnCommand Unified Manager Workflow Guide for Managing Cluster Health

Provides information about using Unified Manager to manage and troubleshoot cluster storage health issues. This guide also describes how to use the Unified Manager maintenance console to perform special operations such as restoring a database backup and connecting to an external data provider to offload performance statistics.

OnCommand Unified Manager Workflow Guide for Managing Cluster Performance

Provides information about using Unified Manager to manage and troubleshoot cluster storage performance issues. This includes identifying workloads that are overusing cluster components so that you can take corrective action to bring performance back to normal levels of operation.

OnCommand Unified Manager Online Help

Provides information about using Unified Manager to manage and troubleshoot cluster storage health and performance issues. Additionally, it provides field level descriptions for every UI page in the product. The online help is included with the software, and is also available as a PDF document that you can review offline.

Common Unified Manager health workflows and tasks

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

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

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

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

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

- Setting up the management environment after deployment
  After storage clusters and their storage resources have been configured using the ONTAP command-line interface (CLI) or System Manager, storage administrators can further specify and configure the clusters for monitoring within Unified Manager.
- Diagnosing and managing availability issues
  If hardware failure or storage resource configuration issues cause the display of data availability events in the Dashboards/Overview page, storage administrators can follow the embedded links to view connectivity information about the affected storage resource, view troubleshooting advice, and assign issue resolution to other administrators.
- Configuring and monitoring performance incidents
  The OnCommand Administrator can monitor and manage the performance of the storage system resources that are being monitored. See the *Unified Manager Workflow Guide for Managing Cluster Performance* for more information.
- Creating, configuring, monitoring, and protecting Infinite Volumes
  After using the Workflow Automation tool to create, configure, and define storage classes for an Infinite Volume, storage administrators can use Unified Manager to monitor, set notification thresholds, and define the data policy for that volume and its storage classes. Optionally, storage administrators can use WFA and Unified Manager to set up data protection for the Infinite Volume.
- Diagnosing and managing volume capacity issues
  If volume storage capacity issues are displayed in the Dashboards/Overview page, storage administrators can follow the embedded links to view the current and historical trends related to the storage capacity of the affected volume, view troubleshooting advice, and assign issue resolution to other administrators.
- Configuring, monitoring, and diagnosing protection relationship issues
  After creating and configuring protection relationships, storage administrators can view the potential issues related to protection relationships in the Dashboards/Overview page, and they can follow the embedded links to view the current state of the protection relationships, the current and historical protection job success information about the affected relationships, and troubleshooting advice, and to assign issue resolution to other administrators. Storage administrators can also configure and manage SnapMirror and SnapVault relationships.

- Creating backup files and restoring data from backup files.
- Associating storage objects with annotations
  By associating storage objects with annotations, storage administrators can filter and view the events that are related to the storage objects, which enables storage administrators to prioritize and resolve the issues that are associated with the events.
- Sending a support bundle to technical support
  Storage administrators can retrieve and send a support bundle to technical support by using the maintenance console. Support bundles must be sent to technical support when the issue requires more detailed diagnosis and troubleshooting than what an AutoSupport message provides.
- Creating new reports for import
  Storage administrators can create new `.rptdesign` files by using the Eclipse plug-in for Business Intelligence and Reporting Tools (BIRT). These reports can be imported to the Unified Manager UI and viewed in the Reports page.
  The reports that are displayed on the Reports page provide the current status of the storage objects. You can make important decisions—such as decisions about storage procurement—based on the current usage. These reports provide a detailed view of storage objects such as volumes, disk shelves, and aggregates.
  The Reports page in the Unified Manager UI 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 from the Reports page. You can also schedule, share, and import a report from this page.

Related concepts

Monitoring and troubleshooting data availability on page 28
Unified Manager monitors the reliability with which authorized users can access your stored data, alerts you to conditions that block or impede that access, and enables you to diagnose those conditions and assign and track their resolution.

Creating, monitoring, and troubleshooting protection relationships on page 38
Unified Manager enables you to create protection relationships, to monitor and troubleshoot mirror protection and backup vault protection of data stored on managed clusters, and to restore data when it is overwritten or lost.

Prioritizing storage object events using annotations on page 70
You can create and apply annotation rules to storage objects so that you can identify and filter those objects based on the type of annotation applied and its priority.

Related tasks

Configuring your environment after deployment on page 10
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.

Setting up and monitoring an SVM with Infinite Volume without storage classes on page 97
You should use OnCommand Workflow Automation (WFA) and Unified Manager to set up and monitor storage virtual machines (SVMs) with Infinite Volume. You should create the SVM with Infinite Volume using WFA and then monitor the Infinite Volume using Unified Manager. Optionally, you can configure data protection for your Infinite Volume.

Managing your Infinite Volume with storage classes and data policies on page 99
You can effectively manage your Infinite Volume by creating the Infinite Volume with the required number of storage classes, configuring thresholds for each storage class, creating rules and a data policy to determine the placement of data written to the Infinite Volume, configuring data protection, and optionally configuring notification alerts.

Resolving capacity issues on page 36
This workflow provides an example of how you can resolve a capacity issue. In this scenario, you are an administrator or operator and you access the Unified Manager Dashboards/Overview page

to see if any of the monitored storage objects have capacity issues. You see that there is a volume with a capacity risk, and you want to determine the possible cause of and resolution to the problem.

Setting up protection relationships in Unified Manager on page 38
There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Restoring data from Snapshot copies on page 54
When you lose data due to a disaster or because directories or files have been accidentally deleted, you can use Unified Manager to locate and restore the data from a Snapshot copy.

Sending a Unified Manager support bundle to technical support on page 105
This workflow shows you how to generate, retrieve, and send a support bundle to technical support using the Unified Manager maintenance console. You should send a support bundle when the issue that you have requires more detailed diagnosis and troubleshooting than an AutoSupport message provides.

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*

  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*

  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.

Related reference

Unified Manager user roles and capabilities on page 163

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

Changing the Unified Manager virtual appliance host name

The network host is assigned a name when the Unified Manager virtual appliance is first deployed. 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

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*

   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. *Restart the Unified Manager virtual machine*

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

Generating an HTTPS security certificate

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.

Before you begin

You must have the OnCommand Administrator role.

About this task

If you do not have access to the Unified Manager web UI, you can regenerate the HTTPS certificate using the maintenance console.

Steps

1. In the toolbar, click , and then click **HTTPS Certificate** from the **Setup** menu.

2. Click **Regenerate HTTPS Certificate**.
   The Regenerate HTTPS Certificate dialog box appears to warn you that the new certificate will replace the existing certificate.
3. If you want to define or change any of the values in the certificate, enter the changes in the **Regenerate HTTPS Certificate** dialog box.

   The Common Name and Alternative Names fields will use the values from the existing certificate if you do not enter new values. The other fields do not require values, but you can enter values, for example, for the City, State, and Country if you want those values to be populated in the certificate.
4. Click **Yes** to regenerate the certificate.
5. Restart the Unified Manager server so that the new certificate takes effect.

After you finish

Verify the new certificate information by viewing the HTTPS certificate.

Related concepts

What a database restore is on page 81
Database restore is the process of copying a backup file from secondary storage to a disk to restore a Unified Manager image to the same or a different system. You perform the restore operation from the Unified Manager console.

Related tasks

Changing the Unified Manager virtual appliance host name on page 11
The network host is assigned a name when the Unified Manager virtual appliance is first deployed. You can change the host name after deployment. If you change the host name, you must also regenerate the HTTPS certificate.

Restarting the Unified Manager virtual machine

You can restart the virtual machine from the maintenance console of Unified Manager. You must restart after generating a new security certificate or if there is a problem with the virtual machine.

Before you begin

The virtual appliance is powered on.

You are logged in to the maintenance console as the maintenance user.

About this task

You can also restart the virtual machine from vSphere by using the **Restart Guest** option. See the VMware documentation for more information.

Steps

1. Access the maintenance console.
2. Select **System Configuration** > **Reboot Virtual Machine**.

Related tasks

Changing the Unified Manager virtual appliance host name on page 11
The network host is assigned a name when the Unified Manager virtual appliance is first deployed. You can change the host name after deployment. If you change the host name, you must also regenerate the HTTPS certificate.

Changing the 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. The new certificate does not take effect until the Red Hat Enterprise Linux machine is restarted.

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.

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 mysqld stop
   ```

3. If you are using a Red Hat 7.x system, change the host name using the Linux `hostnamectl` command:

   ```
   hostnamectl set-hostname new_FQDN

   hostnamectl set-hostname nuhost.corp.widget.com
   ```

4. If you are using a Red Hat 6.x system, change the host name following these steps:
   a. 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

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

   b. 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

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

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

      ```
      hostname new_FQDN

      hostname nuhost.corp.widget.com
      ```

5. Regenerate the HTTPS certificate for the server:

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

6. Restart the network service:

   ```
   service network restart
   ```

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

   ```
   ping new_hostname
   ```

   ```
   ping nuhost
   ```

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

8. 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 ocie start
   ```

   ```
   service ocieau start
   ```

Adding clusters

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

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 the ONTAP system.
  You can verify this by using the following 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.

You can monitor a single cluster by two instances of Unified Manager provided that you have configured a second cluster-management LIF on the cluster so that each instance of Unified Manager connects through a different LIF.

Steps

1. In the left navigation pane, click **Configuration** > **Cluster Data Sources**.
2. From the **Cluster Data Sources** page, click **Add**.
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 and port 443 are 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. Click **Save**.
5. 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 ONTAP.

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

Result

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

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

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 based on the receipt of events.

Steps

1. *Configure notification settings*

   If you want alert notifications sent when certain events occur in your environment, you must configure an SMTP server and supply an email address from which the alert notification will be sent. If you want to use SNMP traps, you can select that option and provide the necessary information.

2. *Enable remote authentication*

   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 users*

   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.

4. *Enable SAML authentication*

   If you want all remote users to be authenticated through a secure Identity provider (IdP) before they can log into the Unified Manager web UI, then you must configure SAML authentication.

5. *Edit global health threshold settings*

   You can modify the health 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.

6. *Add alerts*

   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.

Configuring event notification settings

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

Before you begin

You must have the following information:

- Email address from which the alert notification is sent
  The email address appears in the "From" field in sent alert notifications. If the email cannot be delivered for any reason, this email address is also used as the recipient for undeliverable mail.
- SMTP server host name, and the user name and password to access the server
- SNMP version, trap destination host IP address, outbound trap port, and the community to configure the SNMP trap

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. In the toolbar, click , and then click **Notifications** in the left Setup menu.
2. In the **Setup/Notifications** page, configure the appropriate settings and click **Save**.

   Notes:

   - If the From Address is pre-filled with the address "OnCommand@localhost.com", you should change it to a real, working email address to make sure that all email notifications are delivered successfully.
   - 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

Configuring Unified Manager to send alert notifications on page 15

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.

Enabling remote authentication

You can enable remote authentication so that the Unified Manager server can communicate with your authentication servers. The users of the authentication server can access the Unified Manager graphical interface to manage storage objects and data.

Before you begin

You must have the OnCommand Administrator role.

Important: The Unified Manager server must be connected directly with the authentication server. You must disable any local LDAP clients such as SSSD (System Security Services Daemon) or NSLCD (Name Service LDAP Caching Daemon).

About this task

You can enable remote authentication using either Open LDAP or Active Directory. 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.

Note: The certificate that is used to authenticate users must conform to the X.509 format.

Steps

1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
2. In the **Setup/Authentication** page, select **Enable Remote Authentication**.
3. In the **Authentication Service** field, select the type of service and 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)
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.

4. Optional: Add authentication servers, and test the authentication.

5. Click **Save and Close**.

Related tasks

Configuring Unified Manager to send alert notifications on page 15
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.

Adding authentication servers on page 18
You can add authentication servers and enable remote authentication on the management server so that remote users within the authentication server can access Unified Manager.

Enabling SAML authentication on page 21
You can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

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 have the OnCommand Administrator role.
- Disabling nested groups is only applicable when using Active Directory.

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. In the toolbar, click , and then click **Authentication** in the left Setup menu.
2. In the **Setup/Authentication** page, check the **Disable Nested Group Lookup** box.
3. Click **Save**.

Adding authentication servers

You can add authentication servers and enable remote authentication on the management server so that remote users within the authentication server can 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. In the toolbar, click , and then click **Authentication** in the left Setup menu.
2. In the **Setup/Authentication** page, click **Management Server** > **Authentication**.
3. Required: Enable or disable the **Use secure connection authentication** option:

If you want to...	Then do this...
Enable it	**a.** In Enable remote authentication checkbox, select the **Use Secure Connection** option. **b.** In the Authentication 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.
Disable it	**a.** In the Enable remote authentication checkbox, clear the **Use Secure Connection** option. **b.** In the Authentication 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.

4. Perform a test authentication to confirm that you can authenticate users in the authentication server that you added.

Related tasks

Configuring Unified Manager to send alert notifications on page 15
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.

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. In the toolbar, click , and then click **Authentication** in the left Setup menu.
2. In the **Setup/Authentication** page, click **Test Authentication**.
3. 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.

Adding users

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

Before you begin

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

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

About this task

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

Steps

1. In the toolbar, click , and then click **Users** in the left Management menu.
2. On the **Management/Users** page, click **Add**.
3. In the **Add User** dialog box, select the type of user that you want to add, and enter the required information.

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

Related tasks

Enabling remote authentication on page 17
You can enable remote authentication so that the Unified Manager server can communicate with your authentication servers. The users of the authentication server can access the Unified Manager graphical interface to manage storage objects and data.

Enabling SAML authentication on page 21

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

Related reference

Definitions of user types on page 163
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.

Definitions of user roles on page 162
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 user roles and capabilities on page 163
Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

Enabling SAML authentication

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

Before you begin

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

About this task

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

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

Note: Unified Manager is restarted automatically after SAML authentication is configured.

Steps

1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
2. In the **Setup/Authentication** page, select the **SAML Authentication** tab.
3. Select the **Enable SAML authentication** checkbox.

   The fields required to configure the IdP connection are displayed.
4. Enter the IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP server.

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

5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.
6. Click **Save**.
   A message box displays to confirm that you want to complete the configuration and restart Unified Manager.
7. Click **Confirm and Logout** and Unified Manager is restarted.

After you finish

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

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

Related concepts

What a database restore is on page 81
Database restore is the process of copying a backup file from secondary storage to a disk to restore a Unified Manager image to the same or a different system. You perform the restore operation from the Unified Manager console.

Related tasks

Enabling remote authentication on page 17
You can enable remote authentication so that the Unified Manager server can communicate with your authentication servers. The users of the authentication server can access the Unified Manager graphical interface to manage storage objects and data.

Adding users on page 20
You can add local users or database users by using the Management/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.

Disabling SAML authentication from the maintenance console on page 90
You may need to disable SAML authentication from the maintenance console when there is no access to the Unified Manager GUI. This could happen in cases of mis-configuration or if the IdP is not accessible.

Related reference

Identity provider requirements on page 85
When configuring Unified Manager to use an identity provider (IdP) to perform SAML authentication for all remote users, you need to be aware of some required configuration settings so that the connection to Unified Manager is successful.

Configuring global health threshold settings

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

Global health threshold settings are accessible from the Configuration/Health Thresholds page. You can also modify threshold settings for individual objects, from the inventory page or the details page for that object.

Choices

- *Configuring global aggregate health threshold values*

  You can configure the health threshold settings for capacity, growth, and Snapshot copies for all aggregates to track any threshold breach.
- *Configuring global volume health threshold values*

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

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

Related tasks

Configuring Unified Manager to send alert notifications on page 15

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.

Configuring global aggregate health threshold values

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

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

Steps

1. In the left navigation pane, click **Configuration** > **Health Thresholds**.
2. In the **Configuration/Health Thresholds** page, click **Aggregates**.
3. Configure the appropriate threshold values for capacity, growth, and Snapshot copies.
4. Click **Save**.

Configuring global volume health threshold values

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

Before you begin

You must have the 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.

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

Steps

1. In the left navigation pane, click **Configuration** > **Health Thresholds**.
2. In the **Configuration/Health Thresholds** page, click **Volumes**.
3. Configure the appropriate threshold values for capacity, Snapshot copies, qtree quotas, volume growth, and inodes.
4. Click **Save**.

Editing lag health threshold settings for unmanaged protection relationships

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

Before you begin

You must have the 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. In the left navigation pane, click **Configuration** > **Health Thresholds**.
2. In the **Configuration/Health Thresholds** page, click **Relationships**.
3. Increase or decrease the global default warning or error lag time percentage as required.
4. Click **Save**.

Subscribing to ONTAP EMS events

You can subscribe to receive Event Management System (EMS) events that are generated by systems that are installed with ONTAP software. EMS events are reported to Unified Manager only if you have subscribed to these events.

About this task

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 clusters you are monitoring in Unified Manager.

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.

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 preventing the cluster from reaching the Unified Manager server. To resolve this issue, the cluster administrator must correct the DNS configuration of the cluster, and then restart Unified Manager. Doing so will flush the pending EMS events to the Unified Manager server.

Steps

1. In the left navigation pane, click **Configuration** > **Manage Events**.
2. In the **Configuration/Manage Events** page, click the **Subscribe to EMS events** button.

3. 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 events to which you can subscribe, from the ONTAP cluster shell, you can use the `event route show` command (prior to ONTAP 9) or the `event catalog show` command (ONTAP 9 or later).

   How to configure ONTAP EMS Event Subscriptions in OnCommand Unified Manager / Active IQ Unified Manager

4. Click **Add**.
   The EMS event is added to the Subscribed EMS events list, but the Applicable to Cluster column displays the status as "Unknown" for the EMS event that you added.
5. Click **Save and Close** to register the EMS event subscription with the cluster.
6. Click **Subscribe to EMS events** again.
   The status "Yes" appears in the Applicable to Cluster column for the EMS event that you added.

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

After you finish

When the ONTAP EMS event occurs, the event is displayed on the Events page. You can select the event to view details about the EMS event in the Event details page. You can also manage the disposition of the event or create alerts for the event.

Adding alerts

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

Before you begin

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

About this task

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

Steps

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

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

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

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

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

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

   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 Management/Users page, the modified email address is not updated for the selected user.

   You can also choose to notify users through SNMP traps.

7. Click **Save**.

Example of adding an alert

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

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

Perform the following steps in the Add Alert dialog box:

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

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

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

Related tasks

Configuring Unified Manager to send alert notifications on page 15

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.

Related reference

Description of event severity types on page 114
Each event is associated with a severity type to help you prioritize the events that require immediate corrective action.

Description of event impact levels on page 115
Each event is associated with an impact level (Incident, Risk, or Event) to help you prioritize the events that require immediate corrective action.

Configuring database backup settings

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

Before you begin

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

  It is recommended that you use a remote location that is external to the Unified Manager host system.
- When Unified Manager is installed on Red Hat Enterprise Linux, verify that the "jboss" user has write permissions to the backup directory.
- You should not schedule backup operations to occur immediately after a new cluster has been added while Unified Manager is collecting 15 days of historical performance data.

About this task

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

Steps

1. In the toolbar, click , and then click **Management** > **Database Backup**.
2. In the **Management/Database Backup** page, click **Actions** > **Database Backup Settings**.
3. 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.
4. In the **Schedule Frequency** section, select the **Enable** checkbox, and then 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.
5. Click **Save and Close**.

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**.

Monitoring and troubleshooting data availability

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

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

Related tasks

Resolving a flash card offline condition on page 28
This workflow provides an example of how you might resolve a flash card offline condition. In this scenario, you are an administrator or operator monitoring the dashboard to check for problems with availability. You see a flash card offline condition and you want to determine the possible cause of and resolution to the problem.

Scanning for and resolving storage failover interconnect link down conditions on page 30
This workflow provides an example of how you might scan for, evaluate, and resolve downed storage failover interconnect link conditions. In this scenario, you are an administrator using Unified Manager to scan for storage failover risks before starting an ONTAP version upgrade on your nodes.

Resolving volume offline issues on page 32
This workflow provides an example of how you might evaluate and resolve a volume offline event that Unified Manager might display in the Availability area of the Dashboards/Overview page. In this scenario, you are an administrator using Unified Manager to troubleshoot one or more volume offline events that are displayed on the Dashboards/Overview page.

Resolving a flash card offline condition

This workflow provides an example of how you might resolve a flash card offline condition. In this scenario, you are an administrator or operator monitoring the dashboard to check for problems

with availability. You see a flash card offline condition and you want to determine the possible cause of and resolution to the problem.

Before you begin

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

About this task

The event information and links displayed in the Availability area of the Unified Manager Dashboards/Overview page monitor the overall availability of data storage resources on the monitored clusters enable you to diagnose specific events that might affect that availability.

In this scenario, the Dashboards/Overview page displays the event Flash Cards Offline in its Availability Incidents section. If a flash card is offline, availability of stored data is impeded because the performance of the cluster node on which it is installed is impaired. You can perform the following steps to localize and identify the potential problem:

Steps

1. From the **Availability** panel in the **Unresolved Incidents and Risks** section, click the hypertext link displayed for Flash Cards Offline.
   The Event details page for the availability incident is displayed.
2. On the **Event** details page, you can review the information displayed in the Cause field and perform one or more of the following tasks:
   - Assign the event to an administrator. *Assigning events*
   - Click the source of the event, in this case the cluster node on which the offline flash card is located, to get more information about that node. *Performing corrective action for a flash card offline*
   - Acknowledge the event. *Acknowledging and resolving events*

Performing corrective action for a flash card offline

After reviewing the description in the Cause field of the Flash Card Offline Event details page, you can search for additional information helpful to resolving the condition.

Before you begin

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

About this task

In this example scenario, the event summary provided on the Event details page contains the following information about the offline flash card condition:

```
Severity: Critical
State: New
Impact Level: Incident
Impact Area: Availability
Source: alpha-node
Source Type: Node
Acknowledged By:
Resolved By:
Assigned To:
Cause: Flash cards at slot numbers 3 are offline.
Alert Settings:
```

The event information indicates that the flash card installed in slot 3 in the cluster node named "alpha-node" is offline.

The information localizes the flash card offline condition to a specific slot on a specific cluster node but does not suggest a reason that the flash card is offline.

Steps

1. To obtain further details that might help you diagnose the flash card offline condition, you can click the name of the source of the event.

   In this example, the source of the event is the "alpha-node" cluster node. Clicking that node name displays the HA Details on the Nodes tab of the Health/Cluster details page for the affected cluster. The displayed HA Details displays information about the HA pair to which that node belongs.

   In this example, the relevant information is in the Events summary table on the HA Details. The table specifies the flash card offline event, the time the event was generated, and, again, the cluster node from which this event originated.
2. Using the ONTAP CLI or OnCommand System Manager, access the Event Manager System (EMS) logs for the affected cluster.

   In this example, you use the event name, the event time, and the event source to find the EMS report on this event. The EMS report on the event contains a detailed description of the event and often advice to remedy the condition indicated by the event.

After you finish

After you diagnose the problem, contact the appropriate administrator or operator to complete the manual steps necessary to get the flash card back online.

Related reference

Event details page on page 111
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.

Health/Cluster details page on page 142
The Health/Cluster details page provides detailed information about a selected cluster, 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.

Unified Manager user roles and capabilities on page 163
Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

Scanning for and resolving storage failover interconnect link down conditions

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

Before you begin

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

About this task

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

Steps

1. To check for recent availability events related to storage failover issues, check the Availability Incidents section and the Availability Risks listings on the **Dashboards/Overview** page.

2. To check further for all availability events related to storage failover issues, perform the following steps:
   a. Click the **Availability Incidents** link on the **Dashboards/Overview** page.
      The Events inventory page displays all events on the monitored clusters.
   b. On the **Events** inventory page, select the options **Incident** and **Risk** in the Filter column.
   c. At the top of the **Events** inventory page Names column, click  and enter **`*failover`** in the text box to limit the event to display to storage failover-related events.

      All past events related to storage failover conditions are displayed.

      In this scenario, the Unified Manager displays the event, "Storage Failover Interconnect One or More Links Down" in its Availability Incidents section.
3. If one or more events related to storage failover are displayed either on the **Dashboards/Overview** page or on the **Events** inventory page, perform the following steps:
   a. Click the event title link to display event details for that event.
      In this example, you click the event title "Storage Failover Interconnect One or More Links Down".
      The Event details page for that event is displayed.
   b. On the **Event** details page, you can perform one or more of the following tasks:
      - Review the error message in the Cause field and evaluate the issue. *Performing corrective action for storage failover interconnect links down*
      - Assign the event to an administrator. *Assigning events*
      - Acknowledge the event. *Acknowledging and resolving events*

Related reference

Event details page on page 111
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.

Unified Manager user roles and capabilities on page 163
Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

Performing corrective action for storage failover interconnect links down

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

Before you begin

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

About this task

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

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

Summary

Severity: Warning
State: New
Impact Level: Risk
Impact Area: Availability
Source: aardvark
```

```
Source Type: Node
Acknowledged By:
Resolved By:
Assigned To:
Cause: At least one storage failover interconnected link
       between the nodes aardvark and bonobo is down.
       RDMA interconnect is up (Link0 up, Link1 down)
```

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

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

Steps

1. From the **Event** details page, you can click the link to the event specified in the Source field to obtain further details of other events that might be related to the storage failover interconnection link down condition.
   In this example, the source of the event is the node named aardvark. Clicking that node name displays the HA Details for the affected HA pair, aardvark and bonobo, on the Nodes tab of the Health/Cluster details page, and displays other events that recently occurred on the affected HA pair.
2. Review the **HA Details** for more information relating to the event.
   In this example, the relevant information is in the Events table. The table shows the "Storage Failover Connection One or More Link Down" event, the time the event was generated, and, again, the node from which this event originated.

After you finish

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

Related reference

Event details page on page 111
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.

Health/Cluster details page on page 142
The Health/Cluster details page provides detailed information about a selected cluster, 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.

Unified Manager user roles and capabilities on page 163
Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

Resolving volume offline issues

This workflow provides an example of how you might evaluate and resolve a volume offline event that Unified Manager might display in the Availability area of the Dashboards/Overview page. In this scenario, you are an administrator using Unified Manager to troubleshoot one or more volume offline events that are displayed on the Dashboards/Overview page.

Before you begin

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

About this task

Volumes might be reported offline for several reasons:

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

You can use the Dashboards/Overview page and the Health/Cluster, Health/SVM, and Health/Volume details pages to confirm or eliminate one or more of these possibilities.

Steps

1. From the **Availability** panel in the **Unresolved Incidents and Risks** section, click the hypertext link displayed for the Volume Offline event.
   The Event details page for the availability incident is displayed.
2. On that page, check the notes for any indication that the SVM administrator has taken the volume in question offline.
3. On the **Event** details page, you can review the information for one or more of the following tasks:
   - Review the information displayed in the Cause field for possible diagnostic guidance.
     In this example, the information in the Cause field informs you only that the volume is offline.
   - Check the Notes and Updates area for any indication that the SVM administrator has deliberately taken the volume in question offline.
   - Click the source of the event, in this case the volume that is reported offline, to get more information about that volume. *Performing corrective action for volume offline conditions*
   - Assign the event to an administrator. *Assigning events*
   - Acknowledge the event or, if appropriate, mark it as resolved. *Acknowledging and resolving events*

Performing diagnostic actions for volume offline conditions

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

Before you begin

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

About this task

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

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

Choices

- Click **Health/Volume** details page links to determine if the volume is offline because its host node is down and storage failover to its HA pair partner has failed also.
  See *Determining if a volume offline condition is caused by a down node*.
- Click **Health/Volume** details page links to determine if the volume is offline and its host storage virtual machine (SVM) is stopped because the node hosting the root volume of that SVM is down.

See *Determining if a volume is offline and SVM is stopped because a node is down*.

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

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

Related reference

Unified Manager user roles and capabilities on page 163
Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

Health/Volume details page on page 116
You can use the Health/Volume details page to view detailed information about a selected volume, such as capacity, storage efficiency, configuration, protection, annotation, and events generated. You can also view information about the related objects and related alerts for that volume.

Health/Storage Virtual Machine details page on page 128
You can use the Health/Storage Virtual Machine details page to view detailed information about the selected SVM, such as its health, capacity, configuration, data policies, logical interfaces (LIFs), LUNs, qtrees, and user and user group quotas. You can also view information about the related objects and related alerts for the SVM.

Health/Cluster details page on page 142
The Health/Cluster details page provides detailed information about a selected cluster, 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.

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

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

Before you begin

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

About this task

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

Steps

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

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

2. In the **Related Devices** pane of the **Health/Storage Virtual Machine** details page, locate and click hypertext link displayed under Volumes.
   The Health/Volumes inventory page displays a table of information about all the volumes hosted by the SVM.

3. On the **Health/Volumes** inventory page State column header, click the filter symbol , and then select the option **Offline**.
   Only the SVM volumes that are in offline state are listed.

4. On the **Health/Volumes** inventory page, click the grid symbol , and then select the option **Cluster Nodes**.

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

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

5. On the **Health/Volumes** inventory page, locate the listing for the offline volume and, in its Cluster Node column, click the name of its hosting node.
   The Nodes tab on the Health/Cluster details page displays the state of the HA pair of nodes to which the hosting node belongs. The state of the hosting node and the success of any cluster failover operation is indicated in the display.

After you finish

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

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

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

Before you begin

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

About this task

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

Steps

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

   The Health/Storage Virtual Machine details page displays the "running" or the "stopped" status of the hosting SVM. If the SVM status is running, then the volume offline condition is not caused by the node hosting the root volume of that SVM being down.
3. If the SVM status is stopped, then click **View SVMs** to further identify the cause of the hosting SVM being stopped.
4. On the **Health/Storage Virtual Machines** inventory page SVM column header, click the filter symbol  and then type the name of the stopped SVM.
   The information for that SVM is shown in a table.
5. On the**Health/Storage Virtual Machines** inventory page, click  and then select the option **Root Volume**.
   The Root Volume column is added to the SVM inventory and displays the name of the root volume of the stopped SVM.
6. In the Root Volume column, click the name of the root volume to display the **Health/Storage Virtual Machine** details page for that volume.
   If the status of the SVM root volume is (Online), then the original volume offline condition is not caused because the node hosting the root volume of that SVM is down.
7. If the status of the SVM root volume is (Offline), then locate and click the hypertext link displayed under Aggregate in the **Related Devices** pane of the SVM root volume's **Health/Volume** details page.
8. Locate and click the hypertext link displayed under Node in the **Related Devices** pane of the Aggregate's **Health/Aggregate** details page.

The Nodes tab on the Health/Cluster details page displays the state of the HA pair of nodes to which the SVM root volume's hosting node belongs. The state of the node is indicated in the display.

After you finish

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

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

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

Before you begin

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

About this task

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

Steps

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

   The Health/Aggregate details page displays the online or offline status of the hosting aggregate. If the aggregate status is online, then RAID disk problems are not the cause of the volume being offline.
2. If the aggregate status is offline, then click **Disk Information** and look for broken disk events in the **Events** list on the **Disk Information** tab.
3. To further identify the broken disks, click the hypertext link displayed under Cluster in the **Related Devices** pane.
   The Health/Cluster details page is displayed.
4. Click **Disks**, and then select **Broken** in the **Filters** pane to list all disks in the broken state.

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

After you finish

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

Resolving capacity issues

This workflow provides an example of how you can resolve a capacity issue. In this scenario, you are an administrator or operator and you access the Unified Manager Dashboards/Overview page to see if any of the monitored storage objects have capacity issues. You see that there is a volume with a capacity risk, and you want to determine the possible cause of and resolution to the problem.

Before you begin

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

About this task

On the Dashboards/Overview page, you look at the Unresolved Incidents and Risks area and see a "Volume Space Full" error event in the Capacity pane under SVM Volume Capacity at Risk.

Steps

1. In the **Unresolved Incidents and Risks** area of the **Dashboards/Overview** page, click the name of the Volume Space Full error event in the **Capacity** pane.
   The Event details page for the error is displayed.
2. From the **Event** details page, you can perform one or more of the following tasks:
   - Review the error message in the Cause field and click the suggestions under Suggested Remedial Actions to review descriptions of possible remediations. *Performing suggested remedial actions for a full volume*
   - Click the object name, in this case a volume, in the Source field to get details about the object. *Volume details page*
   - Look for notes that might have been added about this event. *Adding and reviewing notes associated with an event*
   - Add a note to the event. *Adding and reviewing notes associated with an event*
   - Assign the event to another user. *Assigning events*
   - Acknowledge the event. *Acknowledging and resolving events*
   - Mark the event as resolved. *Acknowledging and resolving events*

Related reference

Event details page on page 111
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.

Performing suggested remedial actions for a full volume

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

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

About this task

In this example, you have seen a Volume Space Full error event on the Unified Manager Dashboards/Overview page and have clicked the name of the event.

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

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

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

Steps

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

   You decide to resize the volume.

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

Related reference

Event details page on page 111
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.

Health/Volume details page on page 116
You can use the Health/Volume details page to view detailed information about a selected volume, such as capacity, storage efficiency, configuration, protection, annotation, and events generated. You can also view information about the related objects and related alerts for that volume.

Health/Aggregate details page on page 153
You can use the Health/Aggregate details page to view detailed information about the selected aggregate, such as the capacity, disk information, configuration details, and events generated. You can also view information about the related objects and related alerts for that aggregate.

Creating, monitoring, and troubleshooting protection relationships

Unified Manager enables you to create protection relationships, to monitor and troubleshoot mirror protection and backup vault protection of data stored on managed clusters, and to restore data when it is overwritten or lost.

Note: Protection operations are not available for FlexGroups.

Related tasks

Resolving a protection job failure on page 49
This workflow provides an example of how you might identify and resolve a protection job failure from the Unified Manager dashboard.

Resolving lag issues on page 52
This workflow provides an example of how you might resolve a lag issue. In this scenario, you are an administrator or operator accessing the Unified Manager Dashboards/Overview page to see if there are any problems with your protection relationships and, if they exist, to find solutions.

Setting up protection relationships in Unified Manager on page 38
There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Performing a protection relationship failover and failback on page 45
When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Setting up protection relationships in Unified Manager

There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Before you begin

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

- You must have established peer relationships between two clusters or two storage virtual machines (SVMs).
- OnCommand Workflow Automation must be integrated with Unified Manager:
  - *Set up OnCommand Workflow Automation*
  - *Verifying Unified Manager data source caching in Workflow Automation*

Steps

1. Depending on the type of protection relationship you want to create, do one of the following:
   - *Create a SnapMirror protection relationship*.
   - *Create a SnapVault protection relationship*.
2. If you want to create a policy for the relationship, depending on the relationship type you are creating, do one of the following:
   - *Create a SnapVault policy*.
   - *Create a SnapMirror policy*.
3. *Create a SnapMirror or SnapVault schedule*.

Configuring a connection between Workflow Automation and Unified Manager

You can configure a secure connection between OnCommand 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

- The installed version of Workflow Automation must be 4.1 or greater.
- You must have installed "WFA pack for managing Clustered Data ONTAP" version 1.4.x on the WFA server. You can download the required pack from the NetApp Storage Automation Store.
  WFA pack for managing Clustered Data ONTAP
- 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. In the toolbar, click , and then click **Workflow Automation** in the left Setup menu.
2. In the **OnCommand Unified Manager Database User** area of the **Setup/Workflow Automation** page, select the name, and enter the password for the database user that you created to support Unified Manager and Workflow Automation connections.
3. In the **OnCommand Workflow Automation Credentials** area of the **Setup/Workflow Automation** page, enter 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).
4. Click **Save** .
5. If you use a self-signed certificate, click **Yes** to authorize the security certificate.
   The Setup/Workflow Automation page displays.

6. Click **Yes** to reload the web UI, and add the Workflow Automation features.

Related tasks

Setting up protection relationships in Unified Manager on page 38
There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Creating a database user on page 95
To support a connection between Workflow Automation and Unified Manager, or to access database views, you must first create a database user with the Integration Schema or Report Schema role in the Unified Manager web UI.

Performing a protection relationship failover and failback on page 45
When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Related information

NetApp Documentation: OnCommand Workflow Automation (current releases)

Verifying Unified Manager data source caching in Workflow Automation

You can determine whether Unified Manager data source caching is working correctly by checking if data source acquisition is successful in Workflow Automation. You might do this when you integrate Workflow Automation with Unified Manager to ensure that Workflow Automation functionality is available after the integration.

Before you begin

You must be assigned either the Administrator role or the Architect role in Workflow Automation to perform this task.

Steps

1. From the Workflow Automation UI, select **Execution** > **Data Sources**.
2. Right-click the name of the Unified Manager data source, and then select **Acquire Now**.
3. Verify that the acquisition succeeds without errors.

   Acquisition errors must be resolved for Workflow Automation integration with Unified Manager to succeed.

Creating a SnapMirror protection relationship from the Health/Volume details page

You can use the Health/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.

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 **Health/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 Health/Volume details page.
7. Click the protection configuration job link at the top of the **Health/Volume** details page.
   The job's tasks and details are displayed in the Protection/Job details page.
8. In the **Protection/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 **Health/Volume** details page.
   The new relationship is displayed in the Health/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 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 ONTAP (version 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 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 tasks

Setting up protection relationships in Unified Manager on page 38
There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Configuring a connection between Workflow Automation and Unified Manager on page 39
You can configure a secure connection between OnCommand 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.

Creating a SnapVault protection relationship from the Health/Volume details page

You can create a SnapVault relationship using the Health/Volume details page so that data backups are enabled for protection purposes on volumes.

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 **Health/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 Health/Volume details page.
7. Click the protection configuration job link at the top of the **Health/Volume** details page.
   The Protection/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 Health/Volume details page topology view.

Related tasks

Setting up protection relationships in Unified Manager on page 38
There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Configuring a connection between Workflow Automation and Unified Manager on page 39
You can configure a secure connection between OnCommand 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.

Creating a SnapVault policy to maximize transfer efficiency

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.
- 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 tasks

Setting up protection relationships in Unified Manager on page 38
There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Configuring a connection between Workflow Automation and Unified Manager on page 39
You can configure a secure connection between OnCommand 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.

Creating a SnapMirror policy to maximize transfer efficiency

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.
- 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 tasks

Setting up protection relationships in Unified Manager on page 38
There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Configuring a connection between Workflow Automation and Unified Manager on page 39
You can configure a secure connection between OnCommand 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.

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 tasks

Setting up protection relationships in Unified Manager on page 38
There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Configuring a connection between Workflow Automation and Unified Manager on page 39
You can configure a secure connection between OnCommand 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.

Performing a protection relationship failover and failback

When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation to perform this operation.

Steps

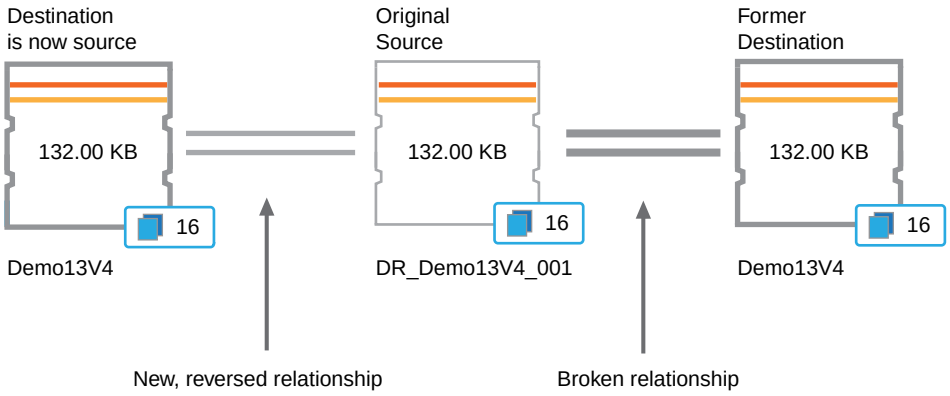
1. *Break the SnapMirror relationship*.

   You must break the relationship before you can convert the destination from a data protection volume to a read/write volume, and before you can reverse the relationship.

2. *Reverse the protection relationship*.

   When the original source volume is available again, you might decide to reestablish the original protection relationship by restoring the source volume. Before you can restore the source, you must synchronize it with the data written to the former destination. You use the reverse resync operation to create a new protection relationship by reversing the roles of the original relationship and synchronizing the source volume with the former destination. A new baseline Snapshot copy is created for the new relationship.

   The reversed relationship looks similar to a cascaded relationship:



3. *Break the reversed SnapMirror relationship*.

   When the original source volume is resynchronized and can again serve data, use the break operation to break the reversed relationship.

4. *Remove the relationship*.

   When the reversed relationship is no longer required, you should remove that relationship before reestablishing the original relationship.

5. *Resynchronize the relationship*.

   Use the resynchronize operation to synchronize data from the source to the destination and to reestablish the original relationship.

Breaking a SnapMirror relationship from the Health/Volume details page

You can break a protection relationship from the Health/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.

Steps

1. In the **Protection** tab of the **Health/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.

Related tasks

Performing a protection relationship failover and failback on page 45
When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Configuring a connection between Workflow Automation and Unified Manager on page 39
You can configure a secure connection between OnCommand 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.

Reversing protection relationships from the Health/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.
- 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 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 **Health/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 Health/Volume details page.
4. Optional: Click **View Jobs** on the **Health/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 **Health/Volume** details page.
   The reverse resynchronization operation is finished when all job tasks are completed successfully.

Related tasks

Performing a protection relationship failover and failback on page 45
When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Configuring a connection between Workflow Automation and Unified Manager on page 39
You can configure a secure connection between OnCommand 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.

Removing a protection relationship from the Health/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.

Steps

1. In the **Protection** tab of the **Health/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 Health/Volume details page.

Resynchronizing protection relationships from the Health/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 have set up OnCommand Workflow Automation.

Steps

1. From the **Protection** tab of the **Health/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 Health/Volume details page and a jobs link is displayed at the top of the page.
9. Optional: Click **View Jobs** on the **Health/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 **Health/Volume** details page.
    The resynchronization job is finished when all job tasks successfully complete.

Related tasks

Performing a protection relationship failover and failback on page 45
When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Configuring a connection between Workflow Automation and Unified Manager on page 39
You can configure a secure connection between OnCommand 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.

Resolving a protection job failure

This workflow provides an example of how you might identify and resolve a protection job failure from the Unified Manager dashboard.

Before you begin

Because some tasks in this workflow require that you log in using the OnCommand Administrator role, you must be familiar with the roles required to use various functionality, as described in *Unified Manager user roles and capabilities* on page 163.

About this task

In this scenario, you access the Dashboards/Overview page to see if there are any issues with your protection jobs. In the Protection Incident area, you notice that there is a Job Terminated incident, showing a Protection Job Failed error on a volume. You investigate this error to determine the possible cause and potential resolution.

Steps

1. In the **Protection Incidents** panel of the Dashboard **Unresolved Incidents and Risks** area, you click the **Protection job failed** event.

   Tip: The linked text for the event is written in the form *`object_name:/object_name - Error Name`*, such as `cluster2_src_svm:/cluster2_src_vol2 - Protection Job Failed`.

   The Event details page for the failed protection job displays.
2. Review the error message in the Cause field of the **Summary** area to determine the problem and evaluate potential corrective actions.

   See *Identifying the problem and performing corrective actions for a failed protection job*.

Related reference

Unified Manager user roles and capabilities on page 163
Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

Identifying the problem and performing corrective actions for a failed protection job

You review the job failure error message in the Cause field on the Event details page and determine that the job failed because of a Snapshot copy error. You then proceed to the Health/ Volume details page page to gather more information.

Before you begin

You must have the OnCommand Administrator role.

About this task

The error message provided in the Cause field on the Event details page contains the following text about the failed job:

```
Protection Job Failed. Reason: (Transfer operation for
relationship 'cluster2_src_svm:cluster2_src_vol2->cluster3_dst_svm:
managed_svc2_vol3' ended unsuccessfully. Last error reported by
Data ONTAP: Failed to create Snapshot copy 0426cluster2_src_vol2snap
on volume cluster2_src_svm:cluster2_src_vol2. (CSM: An operation
failed due to an ONC RPC failure.).)
Job Details
```

This message provides the following information:

- A backup or mirror job did not complete successfully.
  The job involved a protection relationship between the source volume `cluster2_src_vol2` on the virtual server `cluster2_src_svm` and the destination volume `managed_svc2_vol3` on the virtual server named `cluster3_dst_svm`.
- A Snapshot copy job failed for `0426cluster2_src_vol2snap` on the source volume `cluster2_src_svm:/cluster2_src_vol2`.

In this scenario, you can identify the cause and potential corrective actions of the job failure. However, resolving the failure requires that you access either the System Manager web UI or the ONTAP CLI commands.

Steps

1. You review the error message and determine that a Snapshot copy job failed on the source volume, indicating that there is probably a problem with your source volume.

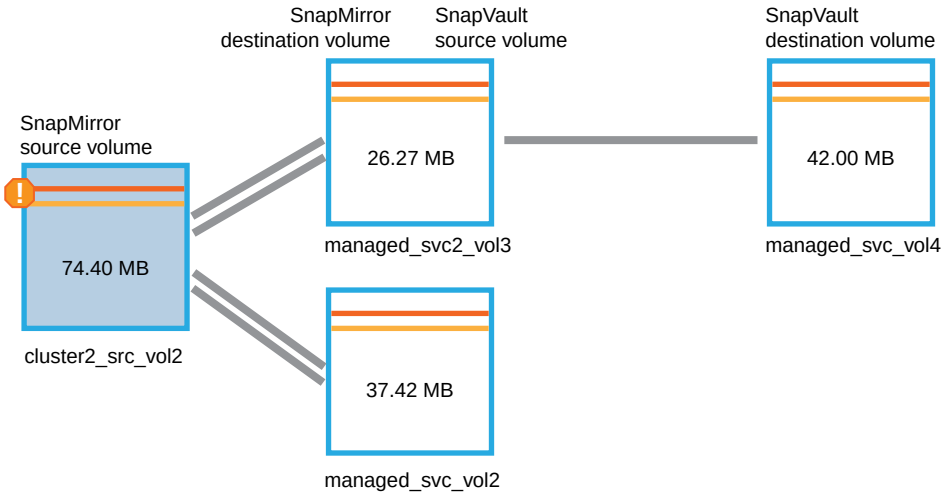
   Optionally, you could click the **Job Details** link at the end of the error message, but for the purposes of this scenario, you choose not to do that.

2. You decide that you want to try to resolve the event, so you do the following:
   a. Click the **Assign To** button and select **Me** from the menu.
   b. Click the **Acknowledge** button so that you do not continue to receive repeat alert notifications, if alerts were set for the event.
   c. Optionally, you can also add notes about the event.
3. Click the **Source** field in the **Summary** pane to see details about the source volume.

   The **Source** field contains the name of the source object: in this case, the volume on which the Snapshot copy job was scheduled.

   The Health/Volume details page displays for `cluster2_src_vol2`, showing the content of the Protection tab.
4. Looking at the protection topology graph, you see an error icon associated with the first volume in the topology, which is the source volume for the SnapMirror relationship.

   You also see the horizontal bars in the source volume icon, indicating the warning and error thresholds set for that volume.



5. You place your cursor over the error icon to see the pop-up dialog box that displays the threshold settings and see that the volume has exceeded the error threshold, indicating a capacity issue.
6. Click the **Capacity** tab.

Capacity information about volume `cluster2_src_vol2` displays.

7. In the **Capacity** pane, you see that there is an error icon in the bar graph, again indicating that the volume capacity has surpassed the threshold level set for the volume.
8. Below the capacity graph, you see that volume autogrow has been disabled and that a volume space guarantee has been set.

   You could decide to enable autogrow, but for the purposes of this scenario, you decide to investigate further before making a decision about how to resolve the capacity problem.
9. You scroll down to the **Events** list and see that Protection Job Failed, Volume Days Until Full, and Volume Space Full events were generated.
10. In the **Events** list, you click the **Volume Space Full** event to get more information, having decided that this event seems most relevant to your capacity issue.
    The Event details page displays the Volume Space Full event for the source volume.
11. In the **Summary** area, you read the Cause field for the event: `The full threshold set at 90% is breached. 45.38 MB (95.54%) of 47.50 MB is used.`
12. Below the **Summary** area, you see Suggested Corrective Actions.

    Tip: The Suggested Corrective Actions display only for some events, so you do not see this area for all types of events.

    You click through the list of suggested actions that you might perform to resolve the Volume Space Full event:

    - Enable autogrow on this volume.
    - Resize the volume.
    - Enable and run deduplication on this volume.
    - Enable and run compression on this volume.
13. You decide to enable autogrow on the volume, but to do so, you must determine the available free space on the parent aggregate and the current volume growth rate:
    a. Look at the parent aggregate, `cluster2_src_aggr1`, in the **Related Devices** pane.

       Tip: You can click the name of the aggregate to get further details about the aggregate.

       You determine that the aggregate has sufficient space to enable volume autogrow.
    b. At the top of the page, look at the icon indicating a critical incident and review the text below the icon.
       You determine that "Days to Full: Less than a day | Daily Growth Rate: 5.4%".
14. Go to System Manager or access the ONTAP CLI to enable the `volume autogrow` option.

    Tip: Make note of the names of the volume and aggregate so you have them available when enabling autogrow.
15. After resolving the capacity issue, return to the Unified Manager **Event** details page and mark the event as resolved.

Related tasks

Adding and reviewing notes about an event on page 110
While addressing events, you can add information about how the issue is being addressed by using the Notes and Updates area in the Event details page. This information can enable another user who is assigned to address the event. You can also view information that was added by the user who last addressed an event, based on the recent timestamp.

Assigning events to specific users on page 110

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.

Acknowledging and resolving events on page 110
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.

Related reference

Protection/Job details page on page 161
The Protection/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.

Resolving lag issues

This workflow provides an example of how you might resolve a lag issue. In this scenario, you are an administrator or operator accessing the Unified Manager Dashboards/Overview page to see if there are any problems with your protection relationships and, if they exist, to find solutions.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

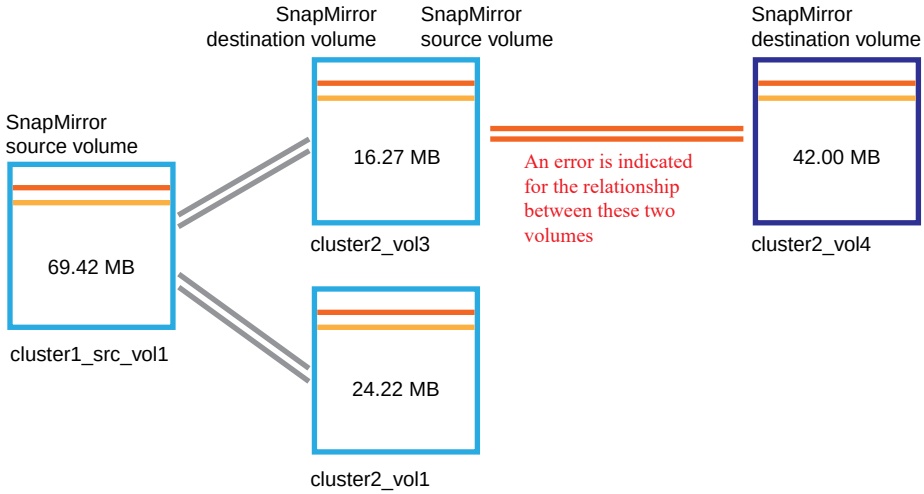
About this task

In the Dashboards/Overview page, you look at the Unresolved Incidents and Risks area and see a SnapMirror Lag error in the Protection pane under Protection Risks.

Steps

1. In the **Protection** pane on the **Dashboards/Overview** page, locate the SnapMirror relationship lag error and click it.
   The Event details page for the lag error event is displayed.
2. From the **Event** details page you can perform one or more of the following tasks:
   - Review the error message in the Cause field of the Summary area to determine if there is any suggested corrective action.
   - Click the object name, in this case a volume, in the Source field of the Summary area to get details about the volume.
   - Look for notes that might have been added about this event.
   - Add a note to the event.
   - Assign the event to a specific user.
   - Acknowledge or resolve the event.
3. In this scenario, you click the object name (in this case, a volume) in the Source field of the **Summary** area to get details about the volume.
   The Protection tab of the Health/Volume details page is displayed.
4. In the **Protection** tab, you look at the topology diagram.

   You note that the volume with the lag error is the last volume in a three-volume SnapMirror cascade. The volume you selected is outlined in dark gray, and a double orange line from the source volume indicates a SnapMirror relationship error.



5. Click each of the volumes in the SnapMirror cascade.
   As you select each volume, the protection information in the Summary, Topology, History, Events, Related Devices, and Related Alerts areas changes to display details relevant to the selected volume.
6. You look at the **Summary** area and position your cursor over the information icon in the **Update Schedule** field for each volume.

   In this scenario, you note that the SnapMirror policy is DPDefault, and the SnapMirror schedule updates hourly at five minutes after the hour. You realize that all of the volumes in the relationship are attempting to complete a SnapMirror transfer at the same time.
7. To resolve the lag issue, you modify the schedules for two of the cascaded volumes so that each destination begins a SnapMirror transfer after its source has completed a transfer.

Related tasks

Adding and reviewing notes about an event on page 110
While addressing events, you can add information about how the issue is being addressed by using the Notes and Updates area in the Event details page. This information can enable another user who is assigned to address the event. You can also view information that was added by the user who last addressed an event, based on the recent timestamp.

Assigning events to specific users on page 110
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.

Acknowledging and resolving events on page 110
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.

Related reference

Event details page on page 111
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.

Unified Manager user roles and capabilities on page 163

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

Restoring data from Snapshot copies

When you lose data due to a disaster or because directories or files have been accidentally deleted, you can use Unified Manager to locate and restore the data from a Snapshot copy.

About this task

You can restore data from two locations in the Unified Manager web UI.

Step

Restore data using one of the following tasks:

- *Restore data from the Health/Volume details page*.
- *Restore data from the Health/Volumes page* .

Restoring data using the Health/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 Health/Volume details page. Restoring data from volumes that have a version of 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 **Health/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.

   Note: If a restore operation fails between ONTAP Cloud HA clusters with an NDMP error, you may need to add an explicit AWS route in the destination cluster so that the destination can communicate with the source system's cluster management LIF. You perform this configuration step using OnCommand Cloud Manager.

Restoring data using the Health/Volumes inventory page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Health/Volumes inventory page. Restoring data from volumes that have a version of 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 **Health/Volumes** inventory 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.

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.

Related concepts

How scripts work with alerts on page 56

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.

Related tasks

Adding scripts on page 57
You can add scripts in Unified Manager, and associate the scripts with alerts. These scripts are executed automatically when an alert is generated, and enable you to obtain information about storage objects for which the event is generated.

Deleting scripts on page 57
You can delete a script from Unified Manager when the script is no longer required or valid.

Testing script execution on page 58
You can verify that your script is executed correctly when an alert is generated for a storage object.

How scripts work with alerts

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

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

The script uses the following arguments for execution:

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

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

Example for obtaining arguments from scripts

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

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

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

Related concepts

Managing scripts on page 55

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.

Adding scripts

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

Before you begin

- You must have created and saved the scripts that you want to add to Unified Manager.
- The supported file formats for scripts are Perl, Shell, PowerShell, and `.bat` files.
  - For Perl scripts, Perl must be installed on the Unified Manager server. If Perl was installed after Unified Manager, you must restart the Unified Manager server.
  - For PowerShell scripts, the appropriate PowerShell execution policy must be set on the server so that the scripts can be executed.
- 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. In the toolbar, click , and then click **Scripts** in the left Management menu.
2. In the **Management/Scripts** page , click **Add**.
3. In the **Add Script** dialog box, click **Browse** to select your script file.
4. Enter a description for the script that you select.
5. Click **Add**.

Related concepts

Managing scripts on page 55
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.

Related tasks

Testing script execution on page 58
You can verify that your script is executed correctly when an alert is generated for a storage object.

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. In the toolbar, click , and then click **Scripts** in the left Management menu.
2. In the **Management/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 concepts

Managing scripts on page 55

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.

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. In the toolbar, click , and then click **Scripts** in the left Management menu.
2. In the **Management/Scripts** page page, add your test script.
3. In the **Configuration/Alerting** page, perform one of the following actions:

To...	Do this...
Add an alert	**a.** In the Configuration/Alerting page, click **Add**. **b.** In the Actions section, associate the alert with your test script.
Edit an alert	**a.** In the Configuration/Alerting 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 **Configuration/Alerting** 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 concepts

Managing scripts on page 55

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.

Related tasks

Adding scripts on page 57

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

Managing and monitoring groups

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

Related concepts

Understanding groups on page 59

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.

How group rules work for groups on page 60

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.

How group actions work on storage objects on page 62
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.

Related tasks

Adding groups on page 63
You can create groups to combine clusters, volumes, and storage virtual machines (SVMs) for ease of management.

Editing groups on page 63
You can edit the name and description of a group that you created in Unified Manager.

Deleting groups on page 64
You can delete a group from Unified Manager when the group is no longer required.

Adding group rules on page 64
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.

Editing group rules on page 66
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.

Deleting group rules on page 66
You can delete a group rule from OnCommand Unified Manager when the group rule is no longer required.

Adding group actions on page 67
You can configure group actions that you want to apply to storage objects in a group. Configuring actions for a group enables you to save time, because you do not have to add these actions to each object individually.

Editing group actions on page 67
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.

Configuring volume health thresholds for groups on page 68
You can configure group-level volume health thresholds for capacity, Snapshot copies, qtree quotas, growth, and inodes.

Deleting group actions on page 69
You can delete a group action from Unified Manager when the group action is no longer required.

Reordering group actions on page 69
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.

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.

Related concepts

Managing and monitoring groups on page 58

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

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

Related tasks

Adding groups on page 63
You can create groups to combine clusters, volumes, and storage virtual machines (SVMs) for ease of management.

Adding group rules on page 64
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.

Adding group actions on page 67
You can configure group actions that you want to apply to storage objects in a group. Configuring actions for a group enables you to save time, because you do not have to add these actions to each object individually.

How group rules work for groups

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

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

Conditions

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

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

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

Storage object type	Applicable operands
Volume	• 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.

- Operand
  The list of operands in Unified Manager changes based on the selected object type. The list includes the object name, owning cluster name, owning SVM name, and annotations that you define in Unified Manager.
- Operator
  The list of operators changes based on the selected operand for a condition. The operators supported in Unified Manager are "Is" and "Contains".
  When you select the "Is" operator, the condition is evaluated for exact match of operand value to the value provided for the selected operand.
  When you select the "Contains" operator, the condition is evaluated to meet one of the following criteria:
  - The operand value is an exact match to the value provided for the selected operand
  - The operand value contains the value provided for the selected operand
- Value
  The value field changes based on the operand selected.

Example of a group rule with conditions

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

- Name contains "vol"
- SVM name is "data_svm"

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

Condition groups

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

Example of a group rule with condition groups

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

- Condition group 1
  - Name contains "vol"
  - SVM name is "data_svm"

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

  Condition group 2 selects all volumes that include "vol" in their names and that are annotated with the data-priority annotation value as "critical".

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

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

Related concepts

Managing and monitoring groups on page 58
You can create groups in Unified Manager to manage storage objects.

How group actions work on storage objects

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

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

Example of unique group actions

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

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

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

Related concepts

Managing and monitoring groups on page 58

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

Adding groups

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

Before you begin

You must have the 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. In the toolbar, click , and then click **Management** > **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 concepts

Managing and monitoring groups on page 58

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

What a group is on page 60

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.

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

Related concepts

Managing and monitoring groups on page 58

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

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

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

Related concepts

Managing and monitoring groups on page 58

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

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

To create....	Do this...
A condition	**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

Managing and monitoring groups on page 58
You can create groups in Unified Manager to manage storage objects.

What a group is on page 60
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.

Related tasks

Editing group rules on page 66
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.

Deleting group actions on page 69

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

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

   Note: You cannot change the target object type for a group rule.
4. Click **Save**.

Related concepts

Managing and monitoring groups on page 58
You can create groups in Unified Manager to manage storage objects.

Related tasks

Adding group rules on page 64
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.

Deleting group actions on page 69
You can delete a group action from Unified Manager when the group action is no longer required.

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

Related concepts

Managing and monitoring groups on page 58

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

Adding group actions

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. In the toolbar, click , and then click **Management** > **Groups**.
2. In the **Group Actions** tab, click **Add**.
3. In the **Add Group Action** dialog box, enter a name and description for the action.
4. From the **Group** menu, select a group for which you want to configure the action.
5. From the **Action Type** menu , select an action type.
   The dialog box expands, enabling you to configure the selected action type with required parameters.
6. Enter appropriate values for the required parameters to configure a group action.
7. Click **Add**.

Related concepts

Managing and monitoring groups on page 58
You can create groups in Unified Manager to manage storage objects.

What a group is on page 60
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.

Related tasks

Editing group actions on page 67
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.

Configuring volume health thresholds for groups on page 68
You can configure group-level volume health thresholds for capacity, Snapshot copies, qtree quotas, growth, and inodes.

Reordering group actions on page 69
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.

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

Related concepts

Managing and monitoring groups on page 58
You can create groups in Unified Manager to manage storage objects.

Related tasks

Adding group actions on page 67
You can configure group actions that you want to apply to storage objects in a group. Configuring actions for a group enables you to save time, because you do not have to add these actions to each object individually.

Deleting group actions on page 69
You can delete a group action from Unified Manager when the group action is no longer required.

Configuring volume health thresholds for groups

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

Related concepts

Managing and monitoring groups on page 58
You can create groups in Unified Manager to manage storage objects.

Related tasks

Adding group actions on page 67

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

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 health threshold, global thresholds are applied to the storage objects in that group. Any object-level health thresholds that are set on the storage object are not impacted.

Steps

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

Related concepts

Managing and monitoring groups on page 58
You can create groups in Unified Manager to manage storage objects.

Related tasks

Adding group rules on page 64
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.

Editing group rules on page 66
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.

Editing group actions on page 67
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.

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

Related concepts

Managing and monitoring groups on page 58
You can create groups in Unified Manager to manage storage objects.

Related tasks

Adding group actions on page 67
You can configure group actions that you want to apply to storage objects in a group. Configuring actions for a group enables you to save time, because you do not have to add these actions to each object individually.

Prioritizing storage object events using annotations

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

Related concepts

Understanding more about annotations on page 71
Understanding the concepts about annotations helps you to manage the events related to the storage objects in your environment.

Related tasks

Adding annotations dynamically on page 73
When you create custom annotations, Unified Manager dynamically associates clusters, storage virtual machines (SVMs), and volumes with the annotations by using rules. These rules automatically assign the annotations to storage objects.

Adding values to annotations on page 74
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.

Deleting annotations on page 74
You can delete custom annotations and their values when they are no longer required.

Viewing the annotation list and details on page 75
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.

Deleting values from annotations on page 75
You can delete values associated with custom annotations when that value no longer applies to the annotation.

Creating annotation rules on page 75
You can create annotation rules that Unified Manager uses to dynamically annotate storage objects such as volumes, clusters, or storage virtual machines (SVMs).

Adding annotations manually to individual storage objects on page 77

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

Editing annotation rules on page 77
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.

Configuring conditions for annotation rules on page 78
You can configure one or more conditions to create annotation rules that Unified Manager applies on the storage objects. The storage objects that satisfy the annotation rule are annotated with the value specified in the rule.

Deleting annotation rules on page 79
You can delete annotation rules from OnCommand Unified Manager when the rules are no longer required.

Reordering annotation rules on page 79
You can change the order in which Unified Manager applies annotation rules to storage objects. Annotation rules are applied to storage objects sequentially based on their rank. When you configure an annotation rule, the rank is least. But you can change the rank of the annotation rule depending on your requirements.

Understanding more about annotations

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

What annotations are

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

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

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

How annotation rules work in Unified Manager

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

- You must associate an annotation rule to an annotation.
- You must associate an object type for an annotation rule; only one object type can be associated for an annotation rule.
- Unified Manager adds or removes annotations from storage objects after each monitoring cycle or when a rule is created, edited, deleted, or reordered.
- An annotation rule can have one or more condition groups, and each condition group can have one or more conditions.

- Storage objects can have multiple annotations. An annotation rule for a particular annotation can also use different annotations in the rule conditions to add another annotation to already annotated objects.

Conditions

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

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

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

Storage object type	Applicable operands
Volume	• 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. When you select the "Is" operator, the condition is evaluated for an exact match of the operand value with the value provided for the selected operand. When you select the "Contains" operator, the condition is evaluated to meet one of the following criteria:

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

Example of an annotation rule with conditions

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

- Name contains "vol"
- SVM name is "data_svm"

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

Condition groups

Condition groups are executed using logical OR, and then applied to storage objects. The storage objects must meet the requirements of one of the condition groups to be annotated. The storage

objects that meet the conditions of all the condition groups are annotated. You can use condition groups to increase the scope of storage objects to be annotated.

Example of an annotation rule with condition groups

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

- Condition group 1
  - Name contains "vol"
  - SVM name is "data_svm"

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

  This condition group annotates all volumes that include "vol" in their names and that are annotated with the data-priority annotation value as "critical".

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

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

Description of predefined annotation values

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

Data-priority:Mission critical

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

Data-priority:High

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

Data-priority:Low

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

Adding annotations dynamically

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. In the toolbar, click , and then click **Annotations** in the left Management menu.
2. In the **Annotations** page, 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. Optional: In the **Annotation Values** section, click **Add** to add values to the annotation.
5. Click **Save and Close**.

Related tasks

Adding values to annotations on page 74

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.

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. In the toolbar, click , and then click **Annotations** in the left Management menu.
2. In the **Annotations** page, select the annotation to which you want to add a value and then click **Add** in the **Values** section.
3. In the **Add Annotation Value** dialog box, specify a value for the annotation.
   The value that you specify must be unique for the selected annotation.
4. Click **Add**.

Related tasks

Adding annotations dynamically on page 73

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

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. In the toolbar, click , and then click **Annotations** in the left Management menu.
2. In the **Annotations** tab, select the annotation that you want to delete.
   The details of the selected annotation are displayed.

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

Result

The selected annotation and its value is deleted.

Related tasks

Deleting values from annotations on page 75

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

Viewing the annotation list and details

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

Steps

1. In the toolbar, click , and then click **Annotations** in the left Management menu.
2. In the **Annotations** tab, click the annotation name to view the associated details.

Deleting values from annotations

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

Before you begin

- You must have the 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. In the toolbar, click , and then click **Annotations** in the left Management menu.
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 74

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

Creating annotation rules

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

Before you begin

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

To create...	Do this...
A condition	**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 71

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

Related tasks

Editing annotation rules on page 77
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.

Reordering annotation rules on page 79
You can change the order in which Unified Manager applies annotation rules to storage objects. Annotation rules are applied to storage objects sequentially based on their rank. When you configure an annotation rule, the rank is least. But you can change the rank of the annotation rule depending on your requirements.

Deleting annotation rules on page 79
You can delete annotation rules from OnCommand Unified Manager when the rules are no longer required.

Adding annotations manually to individual storage objects

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

To add annotation to…	Do this…
Clusters	**a.** Click **Health** > **Clusters**. **b.** Select one or more clusters.
Volumes	**a.** Click **Health** > **Volumes**. **b.** Select one or more volumes.
SVMs	**a.** Click **Health** > **SVMs**. **b.** Select one or more SVMs.

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

Editing annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

   You cannot change the target object type for an annotation rule.
4. Click **Save**.

Related tasks

Creating annotation rules on page 75
You can create annotation rules that Unified Manager uses to dynamically annotate storage objects such as volumes, clusters, or storage virtual machines (SVMs).

Deleting annotation rules on page 79
You can delete annotation rules from OnCommand Unified Manager when the rules are no longer required.

Configuring conditions for annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

Example of configuring a condition for an annotation rule

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

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

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

Deleting annotation rules

You can delete annotation rules from 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. In the toolbar, click , and then click **Annotations** in the left Management menu.
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 tasks

Editing annotation rules on page 77

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.

Reordering annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

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

Configuring backup and restore operations

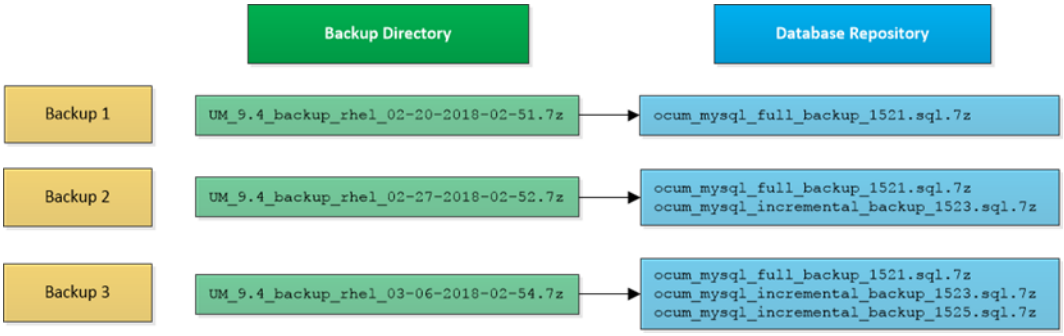
You can create backups of Unified Manager and use the restore feature to restore the backup to the same (local) system or a new (remote) system in case of a system failure or data loss.

What a database backup is

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

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

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



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

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

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

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

Configuring database backup settings

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

Before you begin

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

- You must have a minimum of 150 GB of space available in the location you define as the backup path.

  It is recommended that you use a remote location that is external to the Unified Manager host system.
- When Unified Manager is installed on Red Hat Enterprise Linux, verify that the "jboss" user has write permissions to the backup directory.
- You should not schedule backup operations to occur immediately after a new cluster has been added while Unified Manager is collecting 15 days of historical performance data.

About this task

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

Steps

1. In the toolbar, click , and then click **Management** > **Database Backup**.
2. In the **Management/Database Backup** page, click **Actions** > **Database Backup Settings**.
3. 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.
4. In the **Schedule Frequency** section, select the **Enable** checkbox, and then 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.
5. Click **Save and Close**.

What a database restore is

Database restore is the process of copying a backup file from secondary storage to a disk to restore a Unified Manager image to the same or a different system. You perform the restore operation from the Unified Manager console.

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

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

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

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

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

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

Related tasks

Generating an HTTPS security certificate on page 11
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.

Enabling SAML authentication on page 21
You can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Virtual appliance backup and restore process overview

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

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

1. Power off the VM and take a VMware snapshot of the Unified Manager virtual appliance.
2. Make a NetApp Snapshot copy on the datastore to capture the VMware snapshot.
   If the datastore is not hosted on a system running ONTAP software, follow the storage vendor guidelines to create a backup of the VMware snapshot.
3. Replicate the NetApp Snapshot copy, or snapshot equivalent, to alternate storage.
4. Delete the VMware snapshot.

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

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

Restoring a database backup on a virtual machine

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

Before you begin

- You must have the maintenance user credentials.
- 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, or from a virtual appliance to a Red Hat Enterprise Linux system.

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

Steps

1. In the vSphere client, locate the Unified Manager virtual machine, and then select the **Console** tab.
2. Click in the console window, and then log in to the maintenance console using your user name and password.
3. In the **Main Menu**, enter the number for the **System Configuration** option.
4. In the **System Configuration Menu**, enter the number for the **Restore from an OCUM Backup** option.
5. When prompted, enter the absolute path of the backup file.

```
Bundle to restore from: opt/netapp/data/ocum-backup/
UM_7.2.N151112.0947_backup_unix_11-22-2017-11-41.7z
```

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

After you finish

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

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

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

  It is recommended that you copy the backup file to the default directory `/data/ocum-backup`. The database repository files must be copied to the `/database-dumps-repo` subdirectory under the `/ocum-backup` directory.
- The backup files must be of `.7z` type.

About this task

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

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

Steps

1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
2. Log in as the root user to the Red Hat Enterprise Linux host on which Unified Manager is installed.
3. If Unified Manager is installed in VCS setup, then stop the Unified Manager ocie and ocieau services using Veritas Operations Manager.
4. At the command prompt, restore the backup:

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

   `um backup restore -f /data/ocum-backup/`
   `UM_7.3.N151113.1348_backup_rhel_02-20-2018-04-45.7z`

After you finish

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

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

  It is recommended that you copy the backup file to the default directory `\ProgramData\NetApp\OnCommandAppData\ocum\backup`. The database repository files must be copied to the `\database_dumps_repo` subdirectory under the `\backup` directory.
- The backup files must be of `.7z` type.

About this task

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

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

Steps

1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
2. Log in to the Unified Manager console as an administrator:

   `um cli login -u maint_username`
3. At the command prompt, restore the backup:

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

   `um backup restore -f \ProgramData\NetApp\OnCommandAppData\ocum\backup`
   `\UM_9.4.N151118.2300_backup_windows_02-20-2018-02-51.7z`

After you finish

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

Managing SAML authentication settings

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

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

Identity provider requirements

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

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

Supported encryption standards

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

Validated identity providers

- Shibboleth
- Active Directory Federation Services (ADFS)

ADFS configuration requirements

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

Claim rule	Value
SAM-account-name	Name ID
SAM-account-name	urn:oid:0.9.2342.19200300.100.1.1
Token groups – Unqualified Name	urn:oid:1.3.6.1.4.1.5923.1.5.1.1

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

- Have the CA server reside within the ADFS server to sign the Unified Manager server cert request.

Required Java software

If you are using the Third Party Oracle Java repository with Unified Manager on Windows or Red Hat Enterprise Linux, you must download and install the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files on the Unified Manager server. These files provide for unlimited strength policy files which contain no restrictions on cryptographic strengths.

Note: No change needs to be made when using OpenJDK with Unified Manager.

1. Download the software from .
2. Unzip the two .jar files and copy them to the following location:
   - Red Hat: `$JAVA_HOME/jre/lib/security`
   - Windows: `%JAVA_HOME%\lib\security`

Other configuration requirements

- The Unified Manager clock skew is set to 5 minutes, so the time difference between the IdP server and the Unified Manager server cannot be more than 5 minutes or authentication will fail.
- When users attempt to access Unified Manager using Internet Explorer they might see the message **The website cannot display the page**. If this occurs, make sure these users uncheck the option for "Show friendly HTTP error messages" in **Tools** > **Internet Options** > **Advanced**.

Enabling SAML authentication

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

Before you begin

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

About this task

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

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

Note: Unified Manager is restarted automatically after SAML authentication is configured.

Steps

1. In the toolbar, click , and then click **Authentication** in the left Setup menu.

2. In the **Setup/Authentication** page, select the **SAML Authentication** tab.
3. Select the **Enable SAML authentication** checkbox.

   The fields required to configure the IdP connection are displayed.
4. Enter the IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP server.

   If the IdP server is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URI to populate the IdP Metadata field automatically.
5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.
6. Click **Save**.
   A message box displays to confirm that you want to complete the configuration and restart Unified Manager.
7. Click **Confirm and Logout** and Unified Manager is restarted.

After you finish

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

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

Related concepts

What a database restore is on page 81
Database restore is the process of copying a backup file from secondary storage to a disk to restore a Unified Manager image to the same or a different system. You perform the restore operation from the Unified Manager console.

Related tasks

Enabling remote authentication on page 17
You can enable remote authentication so that the Unified Manager server can communicate with your authentication servers. The users of the authentication server can access the Unified Manager graphical interface to manage storage objects and data.

Adding users on page 20
You can add local users or database users by using the Management/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.

Disabling SAML authentication from the maintenance console on page 90
You may need to disable SAML authentication from the maintenance console when there is no access to the Unified Manager GUI. This could happen in cases of mis-configuration or if the IdP is not accessible.

Related reference

Identity provider requirements on page 85
When configuring Unified Manager to use an identity provider (IdP) to perform SAML authentication for all remote users, you need to be aware of some required configuration settings so that the connection to Unified Manager is successful.

Changing the identity provider used for SAML authentication

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

Before you begin

- You must have the IdP URL and metadata.

- You must have access to the IdP.

About this task

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

Steps

1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
2. In the **Setup/Authentication** page, select the **SAML Authentication** tab.
3. Enter the new IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP.

   If the IdP is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URL to populate the IdP Metadata field automatically.
4. Copy the Unified Manager metadata URI, or save the metadata to an XML text file.
5. Click **Save Configuration**.
   A message box displays to confirm that you want to change the configuration.
6. Click **OK**.

After you finish

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

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

Updating SAML authentication settings after Unified Manager security certificate change

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

About this task

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

Steps

1. Log into the maintenance console.
2. In the **Main Menu**, enter the number for the **Disable SAML authentication** option.
   A message displays to confirm that you want to disable SAML authentication and restart Unified Manager.
3. Launch the Unified Manager user interface using the updated FQDN or IP address, accept the updated server certificate into your browser, and log in using the maintenance user credentials.
4. In the **Setup/Authentication** page, select the **SAML Authentication** tab and configure the IdP connection.
5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.
6. Click **Save**.
   A message box displays to confirm that you want to complete the configuration and restart Unified Manager.
7. Click **Confirm and Logout** and Unified Manager is restarted.

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

Identity provider	Configuration steps
ADFS	**a.** Delete the existing relying party trust entry in the ADFS management GUI. **b.** Add a new relying party trust entry using the `saml_sp_metadata.xml` from the updated Unified Manager server. **c.** Define the three claim rules that are required for Unified Manager to parse ADFS SAML responses for this relying party trust entry. **d.** Restart the ADFS Windows service.
Shibboleth	**a.** Update the new FQDN of Unified Manager server into the `attribute-filter.xml` and `relying-party.xml` files. **b.** Restart the Apache Tomcat web server and wait for port 8005 to come online.

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

Related tasks

Enabling SAML authentication on page 21
You can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Disabling SAML authentication from the maintenance console on page 90
You may need to disable SAML authentication from the maintenance console when there is no access to the Unified Manager GUI. This could happen in cases of mis-configuration or if the IdP is not accessible.

Accessing the maintenance console on page 106
If the Unified Manager user interface is not in operation, or if you need to perform functions that are not available in the user interface, you can access the maintenance console to manage your Unified Manager system.

Disabling SAML authentication

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

About this task

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

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

Note: Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
2. In the **Setup/Authentication** page, select the **SAML Authentication** tab.
3. Uncheck the **Enable SAML authentication** checkbox.

4. Click **Save**.
   A message box displays to confirm that you want to complete the configuration and restart Unified Manager.
5. Click **Confirm and Logout** and Unified Manager is restarted.

Result

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

After you finish

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

Disabling SAML authentication from the maintenance console

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

Before you begin

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

About this task

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

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

Note: Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

1. Log into the maintenance console.
2. In the **Main Menu**, enter the number for the **Disable SAML authentication** option.
   A message displays to confirm that you want to disable SAML authentication and restart Unified Manager.
3. Type **y**, and then press Enter and Unified Manager is restarted.

Result

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

After you finish

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

Related tasks

Enabling SAML authentication on page 21
You can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Accessing the maintenance console on page 106

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

Managing storage objects using the Favorites option

The Favorites option enables you to view and manage selected storage objects in 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 selected storage objects from the Favorites list. The details of a maximum of 20 storage objects are displayed in the Favorites list.

Adding storage objects to the Favorites list

You can 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 when you no longer require them to be marked as favorite.

Adding to, and removing storage objects from, the Favorites list

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

About this task

You can add up to 20 clusters, nodes, aggregates, or volumes to the Favorites list. When you add a node to the Favorites list, it is displayed as a cluster.

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 an aggregate to the Favorites list

1. In the left navigation pane, click **Health** > **Aggregates**.
2. In the Health/Aggregates inventory page, click the aggregate that you want to add to the Favorites list.
3. In the Health/Aggregate details page, click the star icon ().

After you finish

To remove a storage object from the Favorites list, go to the Favorites list page, click the star icon () on the object card you want to remove, and then select the **Remove from Favorites** option.

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.

Position your cursor over the aggregate name to view the name of 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.

Note that for FabricPool aggregates that this information reflects only the capacity on the local performance tier. Click the Capacity tile to view detailed information on the Health/ Aggregate details page.

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.

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.

Note: This section does not appear for FlexGroup volumes.

Performance

Performance statistics 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.

Creating and importing reports into Unified Manager

While Unified Manager provides reporting functionality, you might need to create new reports that are specific to your environment. You can create new reports using the Eclipse Business Intelligence and Reporting Tools (BIRT), and then import them into Unified Manager to view and manage.

Before you begin

You must have the OnCommand Administrator role.

You must have downloaded and installed MySQL Connector/J. You must have the location of the mysql-connector-java-5.1.32-bin.jar file to create the JBDC data source, which connects the report to Unified Manager.

About this task

For more detailed information on creating reports, see the Eclipse BIRT website.

Steps

1. *Download and install MySQL Connector/J* on page 94
   You must download and install the MySQL Connector/J drivers in a specific location. You can use these drivers to create a data source that connects the report to Unified Manager.
2. *Create a database user with the Report Schema role* on page 95
   To support a connection between Workflow Automation and Unified Manager, or to access database views, you must first create a database user with the Integration Schema or Report Schema role in the Unified Manager web UI.
3. *Download the Eclipse plugin for BIRT* on page 95
   To create and import reports to Unified Manager, you must first download the Eclipse Business Intelligence and Reporting Tools (BIRT).
4. *Create a project* on page 95
   Before creating a report for import into Unified Manager, you must first create a project using BIRT.
5. *Create a new report* on page 96
   You can create a new report using the Eclipse plug-in for Business Intelligence and Reporting Tools (BIRT). You might want to create new reports if the existing reports in Unified Manager do not meet the needs of your environment.
6. *Create a JDBC data source* on page 96
   After you have created the new report using BIRT, you must create a data source to connect the report to Unified Manager.
7. *Create a new data set* on page 97
   After creating the data source, you must create a MySQL data set to create the output results for your report. You can also edit the output types after creating the data set.
8. *Import the report into Unified Manager* on page 97
   If you have created a report outside of Unified Manager, you can import and save the report file to use with Unified Manager.

Downloading and installing MySQL Connector/J

You must download and install the MySQL Connector/J drivers in a specific location. You can use these drivers to create a data source that connects the report to Unified Manager.

About this task

You must use MySQL Connector/J version 5.1 or later.

Steps

1. Download the MySQL Connector/J drivers at `dev.mysql.com`.
2. Install the `.jar` file and note its location for future reference.
   For example, install the `.jar` file at `C:\Program Files\MySQL\MySQL Connector J \mysql-connector-java-5.1.32-bin.jar`.

Creating a database user

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

Before you begin

You must have the OnCommand Administrator role.

About this task

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

Steps

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

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

6. Click **Add**.

Downloading the Eclipse Business Intelligence and Reporting Tools (BIRT)

To create and import reports to Unified Manager, you must first download the Eclipse Business Intelligence and Reporting Tools (BIRT).

Step

Download the BIRT software at `http://download.eclipse.org/birt/downloads/`.

After you finish

After downloading the BIRT software, you must extract the resulting .zip file.

Creating a project using BIRT

Before creating a report for import into Unified Manager, you must first create a project using BIRT.

Before you begin

You must have downloaded and extracted the BIRT .zip file.

Steps

1. From the Eclipse interface, select **File** > **New** > **Project**.
2. Expand the **Business Intelligence and Reporting Tools** folder, select **Report Project**, and click **Next**.
3. Type the project name and click **Finish**.

Creating a new report using BIRT

You can create a new report using the Eclipse plug-in for Business Intelligence and Reporting Tools (BIRT). You might want to create new reports if the existing reports in Unified Manager do not meet the needs of your environment.

Before you begin

You must have downloaded and extracted BIRT.

You must have created a project using BIRT.

Steps

1. From the BIRT interface, select **File** > **New** > **Report**.
2. In the **New Report** dialog box and select the project folder, which should be the same as the project folder previously created.

   If you select a different project folder, you cannot use the reporting operations in Unified Manager.
3. Type the report file name, and click **Next**.
4. Select the report type and click **Finish**.

Creating a JDBC data source using BIRT

After you have created the new report using BIRT, you must create a data source to connect the report to Unified Manager.

Before you begin

You must have created a report using BIRT.

You must have downloaded and installed MySQL Connector/J.

You must have created a database user with the Report Schema role.

Steps

1. In Eclipse, select **Data Explorer** > **Data Sources** > **New Data Source**.
2. Select **Create from a data source type in the following list**.
3. Select **JDBC Data Source**, and then click **Next**.
4. In the **New JDBC Data Source Profile** dialog box, select **com.mysql.jdbc.Driver(v5.1)**.
   a. If the MySQL driver does not appear, click **Manage Drivers**.
   b. In the **Manage JDBC Drivers** dialog box, click **Add**.
   c. Browse to the location where the MySQL Connector/J `.jar` file was installed, and then select the file.
   d. Click **OK**.
   You should be able to view and select the MySQL driver.
5. Enter the fully qualified host name or the IP address of the Unified Manager instance using appropriate format:

Address Type	Format
IPv4	`jdbc:mysql://xx.xx.xx.xx:3306/ocum_report`
IPv6	`jdbc:mysql://address=(protocol=tcp)(host=xx:xx:xx:xx:xx:xx:xx:xx)(port=3306)/ocum_report`

6. Enter the user name for the database user, enter the password, and then click **Finish**.

Creating a new MySQL data set using BIRT

After creating the data source, you must create a MySQL data set to create the output results for your report. You can also edit the output types after creating the data set.

Before you begin

You must have created a JDBC data source using BIRT.

You must have downloaded and installed MySQL Connector/J.

You must have created a database user with the Report Schema role in Unified Manager.

Steps

1. From **Eclipse**, select a workspace.
2. Select **Data Explorer** > **Data Sets** > **New Data Set**.
3. In the **New Data Set** dialog box, select the data source previously created, the data set type, and the data set name, and click **Next**.
4. Define an SQL query text using the available items, or manually enter the query, and click **Finish**.
5. Click **Preview Results** to confirm the SQL query, and then click **OK**.
6. In the **Edit Data Set** dialog box, define the output columns as necessary and click **OK**.
7. Drag items into the newly created report.

After you finish

You should now import the newly created report into Unified Manager.

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. In the left navigation pane, click **Reports**, and then click **Import Report**.
2. In the **Import Report** dialog box, click **Browse** and select the file you want to import, and then enter a name and brief description of the report.
3. Click **Import**.

   If you cannot import the report, you can check the log file to find the error causing the issue.

Setting up and monitoring an SVM with Infinite Volume without storage classes

You should use OnCommand Workflow Automation (WFA) and Unified Manager to set up and monitor storage virtual machines (SVMs) with Infinite Volume. You should create the SVM with

Infinite Volume using WFA and then monitor the Infinite Volume using Unified Manager. Optionally, you can configure data protection for your Infinite Volume.

Before you begin

The following requirements must be met:

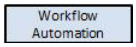
- WFA must be installed and the data sources must be configured.
- You must have the OnCommand Administrator or Storage Administrator role.
- You must have created the required number of aggregates by customizing the appropriate predefined workflow in WFA.
- You must have configured the Unified Manager server as a data source in WFA, and then you must have verified that the data is cached successfully.

About this task

- You can monitor only data SVMs using Unified Manager.
- While performing this task, you are required to switch between two applications: OnCommand Workflow Automation (WFA) and OnCommand Unified Manager.
- The task provides high-level steps.

  For details about performing the WFA tasks, see the *OnCommand Workflow Automation* documentation.

Steps

1. 

Create an SVM with Infinite Volume, and then create the Infinite Volume by using the appropriate workflow.

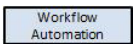
   You can enable storage efficiency technologies, such as deduplication and compression, while creating the Infinite Volume.
2. 

Add the cluster containing the SVM with Infinite Volume to the Unified Manager database.

   You can add the cluster by providing the IP address or the FQDN of the cluster.
3. 

Based on your organization's requirements, modify the thresholds for the Infinite Volume on the SVM.

   Tip: You should use the default Infinite Volume threshold settings.
4. 

Configure notification alerts and traps to address any availability and capacity issues related to the Infinite Volume.
5. Optional: 

Create a disaster recovery (DR) SVM with Infinite Volume, and then configure data protection (DP) by performing the following steps:
   a. Create a data protection (DP) Infinite Volume by using the appropriate workflow.
   b. Set up a DP mirror relationship between the source and destination by using the appropriate workflow.

Related tasks

Editing the Infinite Volume threshold settings on page 99
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.

Adding alerts on page 25

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

Related reference

Unified Manager user roles and capabilities on page 163

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

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. In the left navigation pane, click **Health** > **SVMs**.
2. In the **Health/Storage Virtual Machines** inventory page, select a SVM with Infinite Volume.
3. In the **Health/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**.

Managing your Infinite Volume with storage classes and data policies

You can effectively manage your Infinite Volume by creating the Infinite Volume with the required number of storage classes, configuring thresholds for each storage class, creating rules and a data policy to determine the placement of data written to the Infinite Volume, configuring data protection, and optionally configuring notification alerts.

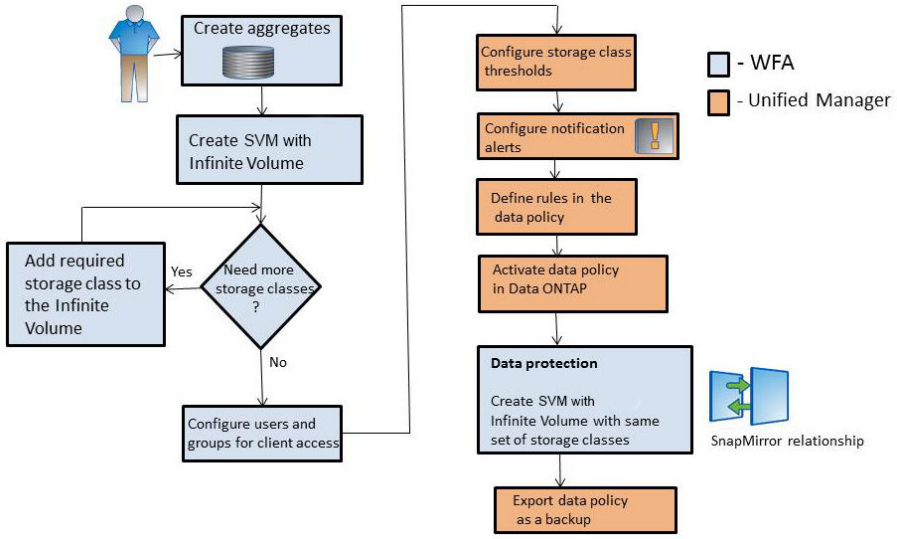
Before you begin

- OnCommand Workflow Automation (WFA) must be installed.
- You must have the OnCommand Administrator or Storage Administrator role.
- You must have created the required number of aggregates by customizing the appropriate predefined workflow in WFA.
- You must have created the required number of storage classes by customizing the appropriate predefined workflow in WFA.
- You must have configured the Unified Manager server as a data source in WFA, and then you must have verified that the data is cached successfully.

About this task

While performing this task, you are required to switch between two applications: OnCommand Workflow Automation (WFA) and OnCommand Unified Manager.

The task provides high-level steps. For details about performing the WFA tasks, see the *OnCommand Workflow Automation* documentation.



Steps

1. Workflow Automation Customize the predefined workflow to define the required storage classes.
2. Workflow Automation Create an SVM with Infinite Volume with the required number of storage classes by using the appropriate workflow.
3. Unified Manager Add the cluster containing the SVM with Infinite Volume to the Unified Manager database.

   You can add the cluster by providing the IP address or the FQDN of the cluster.
4. Unified Manager *Based on your organization's requirements, modify the thresholds for each storage class*.

   You should use the default storage class threshold settings to effectively monitor storage class space.
5. Unified Manager *Configure notification alerts and traps to address any availability and capacity issues related to the Infinite Volume*.
6. Unified Manager *Set up rules in the data policy, and then activate all the changes made to the data policy*

   Rules in a data policy determine the placement of the content written to the Infinite Volume.

   Note: Rules in a data policy affect only new data written to the Infinite Volume and do not affect existing data in the Infinite Volume.
7. Optional: Workflow Automation Create a disaster recovery (DR) SVM with Infinite Volume, and then configure a data protection (DP) by performing the following steps:
   a. Create a data protection (DP) Infinite Volume by using the appropriate workflow.
   b. Set up a DP mirror relationship between the source and destination by using the appropriate workflow.

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. In the left navigation pane, click **Health** > **SVMs**.
2. In the **Health/Storage Virtual Machines** inventory page, select an SVM with Infinite Volume.
3. In the **Health/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**.

Adding alerts

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

Before you begin

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

About this task

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

Steps

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

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

   If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.
5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.

   Tip: To select more than one event, press the Ctrl key while you make your selections.
6. Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.

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 Management/Users page, the modified email address is not updated for the selected user.

You can also choose to notify users through SNMP traps.

7. Click **Save**.

Example of adding an alert

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

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

Perform the following steps in the Add Alert dialog box:

1. Click **Name**, and enter **`HealthTest`** in the **Alert Name** field.
2. Click **Resources**, and in the Include tab, select **Volumes** from the drop-down list.
   a. Enter **`abc`** in the **Name contains** field to display the volumes whose name contains "abc".
   b. Select <<**All Volumes whose name contains 'abc'**>> from the Available Resources area, and move it to the Selected Resources area.
   c. Click **Exclude**, and enter **`xyz`** in the **Name contains** field, and then click **Add**.
3. Click **Events**, and select **Critical** from the Event Severity field.
4. Select **All Critical Events** from the Matching Events area, and move it to the Selected Events area.
5. Click **Actions**, and enter **`sample@domain.com`** in the Alert these users field.
6. Select **Remind every 15 minutes** to notify the user every 15 minutes.
   You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.
7. In the Select Script to Execute menu, select **Test** script .
8. Click **Save**.

Related tasks

Configuring Unified Manager to send alert notifications on page 15
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.

Related reference

Description of event severity types on page 114
Each event is associated with a severity type to help you prioritize the events that require immediate corrective action.

Description of event impact levels on page 115

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

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 103
- *Creating custom rules* on page 104

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 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. In the left navigation pane, click **Health** > **SVMs**.
2. In the **Health/Storage Virtual Machines** inventory 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 the 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.

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. In the left navigation pane, click **Health** > **SVMs**.
2. In the **Health/Storage Virtual Machines** inventory 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.

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. In the left navigation pane, click **Health** > **SVMs**.
2. In the **Health/Storage Virtual Machines** inventory 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.

Sending a Unified Manager support bundle to technical support

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

About this task

For more information about the maintenance console and support bundles, see *Using the maintenance console*.

Unified Manager stores two generated support bundles at one time.

Related concepts

What the maintenance user does on page 170
The maintenance user is created during the installation of Unified Manager. For Red Hat installations the maintenance user name is the "umadmin" user. The maintenance user has the OnCommand administrator role in the web UI, and that user can create subsequent users and assign them roles.

Related reference

Unified Manager user roles and capabilities on page 163
Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

Steps

1. *Access the maintenance console using Secure Shell (SSH)* on page 106
If the Unified Manager user interface is not in operation, or if you need to perform functions that are not available in the user interface, you can access the maintenance console to manage your Unified Manager system.

2. *Generate a support bundle* on page 106
You can generate a support bundle, containing full diagnostic information, so that you can then retrieve it and send it to technical support for troubleshooting help. Because some types of data can use a large amount of cluster resources or take a long time to complete, you can specify data types to include or exclude in the support bundle.

3. *Retrieve the support bundle using a Windows client* on page 108
If you are a Windows user, you can download and install a tool to retrieve the support bundle from your Unified Manager server. You can send the support bundle to technical support for a more detailed diagnosis of an issue. Filezilla or WinSCP are examples of tools you can use.

4. *Retrieve the support bundle using a Unix or Linux client* on page 108
If you are a UNIX or Linux user, you can retrieve the support bundle from your vApp by using the command-line interface (CLI) on your Linux client server. You can use either SCP or SFTP to retrieve the support bundle.

5. *Send the support bundle to technical support* on page 109

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

Accessing the maintenance console

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

Before you begin

You must have installed and configured Unified Manager.

About this task

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

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

Step

Follow these steps to access the maintenance console:

On this operating system...	Follow these steps...
VMware	**a.** Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance. **b.** Log in to the maintenance console using your maintenance user name and password.
Red Hat	**a.** Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system. **b.** Log in to the system with the maintenance user (umadmin) name and password. **c.** Enter the command **`maintenance_console`** and press Enter.
Windows	**a.** Log in to the Unified Manager system with administrator credentials. **b.** Launch PowerShell as a Windows administrator. **c.** Enter the command **`maintenance_console`** and press Enter. **Note:** On Microsoft Windows Server 2012 if you receive an execution policy error, enter the following command and try step c again: **`PowerShell.exe -ExecutionPolicy RemoteSigned`**

The Unified Manager maintenance console menu is displayed.

Related tasks

Using the maintenance console on page 170

You can use the maintenance console to configure network settings, to configure and manage the system on which Unified Manager is installed, and to perform other maintenance tasks that help you prevent and troubleshoot possible issues.

Generating a support bundle

You can generate a support bundle, containing full diagnostic information, so that you can then retrieve it and send it to technical support for troubleshooting help. Because some types of data

can use a large amount of cluster resources or take a long time to complete, you can specify data types to include or exclude in the support bundle.

Before you begin

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

About this task

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

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

Steps

1. In the maintenance console **Main Menu**, select **Support/Diagnostics**.
2. Select **Generate Support Bundle**.
3. Select or deselect the following data types to include or exclude in the support bundle:

   database dump

   A dump of the MySQL Server database.

   heap dump

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

   acquisition recordings

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

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

4. Type **g**, and then press Enter to generate the support bundle.
   Since the generation of a support bundle is a memory intensive operation, you are prompted to verify that you are sure you want to generate the support bundle at this time.
5. Type **y**, and then press Enter to generate the support bundle.

   If you do not want to generate the support bundle at this time, type **n**, and then press Enter.
6. If you included database dump files in the support bundle, you are prompted to specify the time period for which you want database files included:
   a. Enter the starting date in the format YYYYMMDD.

      For example, enter 20170101 for January 1, 2017.
   b. Enter the number of days of statistics to include, beginning from 12 a.m. on the specified starting date.

      You can enter a number from 1 through 10.

   The system displays the period of time for which database statistics will be collected, and then it generates the support bundle.
7. Select **Generate Support Bundle**.

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

After you finish

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

The generated support bundle resides in the `/support` directory on VMware systems, in `/opt/netapp/data/support/` on Red Hat systems, and in `ProgramData\NetApp\OnCommandAppData\ocum\support` on Windows systems.

Related concepts

Diagnostic user capabilities on page 171

The purpose of diagnostic access is to enable technical support to assist you in troubleshooting, and you should only use it when directed by technical support.

Related reference

Unified Manager user roles and capabilities on page 163

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

Retrieving the support bundle using a Windows client

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

Before you begin

You must be the maintenance user to perform this task.

You must use a tool that supports SCP or SFTP.

Steps

1. Download and install a tool to retrieve the support bundle.
2. Open the tool.
3. Connect to your Unified Manager management server over SFTP.
   The tool displays the contents of the `/support` directory and you can view all existing support bundles.
4. Select the destination directory for the support bundle you want to copy.
5. Select the support bundle you want to copy and use the tool to copy the file from the Unified Manager server to your local system.

Related information

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

WinSCP - http://winscp.net

Retrieving the support bundle using a UNIX or Linux client

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

Before you begin

You must be the maintenance user to perform this task.

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

Steps

1. Access the CLI through Telnet or the console, using your Linux client server.
2. Access the `/support` directory.
3. Retrieve the support bundle and copy it to the local directory using the following command:

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

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

4. Enter the maintenance user password.

Examples

The following example uses SCP to retrieve the support bundle:

```
$ scp admin@10.10.12.69:/support/support_bundle_20160216_145359.7z .

Password:
<maintenance_user_password>

support_bundle_20160216_145359.7z   100%  119MB  11.9MB/s   00:10
```

The following example uses SFTP to retrieve the support bundle:

```
$ sftp admin@10.10.12.69:/support/support_bundle_20160216_145359.7z .

Password:
<maintenance_user_password>

Connected to 10.228.212.69.
Fetching /support/support_bundle_20130216_145359.7z to ./support_bundle_20130216_145359.7z
/support/support_bundle_20160216_145359.7z
```

Sending a support bundle to technical support

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

Before you begin

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

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

Steps

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

   How to upload a file to NetApp

Tasks and information related to several workflows

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

Adding and reviewing notes about an event

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

Before you begin

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

Steps

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

Assigning events to specific users

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

Before you begin

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

Steps

1. In the left navigation pane, click **Events**.
2. In the **Events** inventory page, select one or more events that you want to assign.
3. Assign the event by choosing one of the following options:

If you want to assign the event to...	Then do this...
Yourself	Click **Assign To** > **Me**.
Another user	**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.

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.

Note: You cannot acknowledge Information events.

Steps

1. In the left navigation pane, 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**.

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*
- *Summary area*
- *Notes and Updates area*
- *Scope of Impact area*
- *Possible Causes area*
- *Possible Effect area*
- *Resources that Might be Impacted area*
- *Suggested Corrective Actions area*

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.

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

Acknowledge

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

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

Mark As Resolved

Enables you to change the event state to Resolved.

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

Add Alert

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

View Events

Navigates to the Events inventory 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 ().

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

For obsolete events, 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 ONTAP does not provide a valid user name because of SecD errors.

Source Annotations

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

Source Groups

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

Source Type

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

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.

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 area

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

Performing diagnostic actions for volume offline conditions on page 33
After navigating to the Health/Volume details page of a volume reported to be offline, you can search for additional information helpful to diagnosing the volume offline condition.

Performing suggested remedial actions for a full volume on page 37
After receiving a "Volume Space Full" error event, you review the suggested remedial actions on the Event details page and decide to perform one of the suggested actions.

Description of event severity types

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

Critical

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

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

Error

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

Warning

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

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

Information

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

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

Description of event impact levels

Each event is associated with an impact level (Incident, Risk, 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 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.

Protection

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

Health/Volume details page

You can use the Health/Volume details page to view detailed information about a selected volume, such as capacity, storage efficiency, configuration, protection, annotation, and events generated. You can also view information about the related objects and related alerts for that volume.

You must have the OnCommand Administrator or Storage Administrator role.

- *Command buttons*
- *Capacity tab*
- *Efficiency tab*
- *Configuration tab*
- *Protection tab*
- *History area*
- *Events list*
- *Related Annotations pane*
- *Related Devices pane*
- *Related Groups pane*
- *Related Alerts pane*

Command buttons

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

Switch to Performance View

Enables you to navigate to the Performance/Volume details page.



Enables you to add the selected volume to the Favorites dashboard.

Actions

Note that protection actions are not supported for FlexGroup volumes, so those buttons do not appear when viewing the details of a FlexGroup.

- 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

  Note: These buttons are disabled for any version-flexible SnapMirror relationships (XDP) and for all SnapMirror relationships created with ONTAP 9.3. SnapMirror relationships created with ONTAP 9.2, or earlier, and then upgraded to ONTAP 9.3 can use these buttons to actively manage the SnapMirror 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 Health/Volumes inventory 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.

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.
  The number of Snapshot copies in the volume is displayed as a link. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.
  The Snapshot copy count is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

Volume Move

Displays the status of either the current or the last volume move operation that was performed on the volume, and other details, such as the current phase of the volume move operation which is in progress, source aggregate, destination aggregate, start time, end time, and estimated end time.

Also displays the number of volume move operations that are performed on the selected volume. You can view more information about the volume move operations by clicking the **Volume Move History** link.

Efficiency tab

The Efficiency tab displays information about the space saved in the volumes by using storage efficiency features such as deduplication, compression, and FlexClone volumes. Note that the message "Not Applicable" is displayed for efficiency features not supported by FlexGroup volumes.

Deduplication

- Enabled
  Specifies whether deduplication is enabled or disabled on a volume.
- Space Savings
  Displays the amount of space saved (in percentage, or in KB, MB, GB, and so on) in a volume by using deduplication.
- Last Run
  Displays the time that has elapsed since the deduplication operation was last performed. Also specifies whether the deduplication operation was successful.
  If the time elapsed exceeds a week, the timestamp representing when the operation was performed is displayed.
- Mode
  Specifies whether the deduplication operation enabled on a volume is a manual, scheduled, or policy-based operation. If the mode is set to Scheduled, the operation schedule is displayed, and if the mode is set to a policy, the policy name is displayed.
- Status
  Displays the current status of the deduplication operation. The status can be Idle, Initializing, Active, Undoing, Pending, Downgrading, or Disabled.
- Type
  Specifies the type of deduplication operation running on the volume. If the volume is in a SnapVault relationship, the type displayed is SnapVault. For any other volume, the type is displayed as Regular.

Compression

- Enabled
  Specifies whether compression is enabled or disabled on a volume.
- Space Savings

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

Configuration tab

The Configuration tab displays details about the selected volume, such as the export policy, RAID type, capacity and storage efficiency related features of the volume:

Overview

- Full Name
  Displays the full name of the volume.
- Aggregates
  Displays the name of the aggregate on which the volume resides, or the number of aggregates on which the FlexGroup volume resides.
- Tiering Policy
  Displays the tiering policy set for the volume; if the volume is deployed on a FabricPool-enabled aggregate. The policy can be None, Backup, or Snapshot Only.
- Storage Virtual Machine
  Displays the name of the storage virtual machine (SVM) that contains the volume.
- Junction Path
  Displays the status of the path, which can be active or inactive. The path in the SVM to which the volume is mounted is also displayed. You can click the **History** link to view the most recent five changes to the junction path.
- Export policy
  Displays the name of the export policy that is created for the volume. You can click the link to view details about the export policies, authentication protocols, and access enabled on the volumes that belong to the SVM.
- Style
  Displays the volume style. The volume style can be FlexVol or FlexGroup.
- Type
  Displays the type of the selected volume. The volume type can be Read-write, Load-sharing, Data-Protection, Data-cache, or Temporary.
- RAID Type
  Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, or RAID-TEC.

  Note: Multiple RAID types may display for FlexGroup volumes because the constituent volumes for FlexGroups can be on aggregates of different types.
- SnapLock Type
  Displays the SnapLock Type of the aggregate that contains the volume.
- SnapLock Expiry
  Displays the expiry date of SnapLock volume.

Capacity

The message "Not Applicable" is displayed for capacity features not supported by FlexGroup volumes.

- 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.

Note: This tab does not appear for FlexGroup volumes because Unified Manager does not support this functionality.

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 ONTAP software. Yes with backup option indicates the implementation of SnapMirror protection with the ability to retain multiple versions of backup copies on the destination. None indicates that Version Flexible Replication is not enabled.
- Relationship Capability
  Indicates the ONTAP capabilities available to the protection relationship.
- Protection Service
  Displays the name of the protection service if the relationship is managed by a protection partner application.
- Relationship Type
  Displays any relationship type, including 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

Clicking another volume in the topology selects and displays information for that volume.

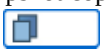
A question mark () in the upper-left corner of a volume indicates that either the volume is missing or that it has not yet been discovered. It might also indicate that the capacity information is missing. Positioning your cursor over the question mark displays additional information, including suggestions for remedial action.

The topology displays information about volume capacity, lag, Snapshot copies, and last successful data transfer if it conforms to one of several common topology templates. If a topology does not conform to one of those templates, information about volume lag and last successful data transfer is displayed in a relationship table under the topology. In that case, the highlighted row in the table indicates the selected volume, and, in the topology view, bold lines with a blue dot indicate the relationship between the selected volume and its source volume.

Topology views include the following information:

- Capacity
  Displays the total amount of capacity used by the volume. Positioning your cursor over a volume in the topology displays the current warning and critical threshold settings for that volume in the Current Threshold Settings dialog box. You can also edit the threshold settings by clicking the **Edit Thresholds** link in the Current Threshold Settings dialog box. Clearing the **Capacity** check box hides all capacity information for all volumes in the topology.
- Lag

Displays the lag duration and the lag status of the incoming protection relationships. Clearing the **Lag** check box hides all lag information for all volumes in the topology. When the **Lag** check box is dimmed, then the lag information for the selected volume is displayed in the relationship table below the topology, as well as the lag information for all related volumes.

- Snapshot
  Displays the number of Snapshot copies available for a volume. Clearing the **Snapshot** check box hides all Snapshot copy information for all volumes in the topology. Clicking a Snapshot copy icon () displays the Snapshot copy list for a volume. The Snapshot copy count displayed next to the icon is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.
- Last Successful Transfer
  Displays the amount, duration, time, and date of the last successful data transfer. When the **Last Successful Transfer** check box is dimmed, then the last successful transfer information for the selected volume is displayed in the relationship table below the topology, as well as the last successful transfer information for all related volumes.

History

Displays in a graph the history of incoming SnapMirror and SnapVault protection relationships for the selected volume. There are three history graphs available: incoming relationship 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. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Graphs might be empty and the message `No data found` displayed when the data or the state of the volume remains unchanged for a period of time.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends—for example, if the volume usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Volume Capacity Used

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

Volume Capacity Used vs Total

Displays the trend in how volume capacity is used based on the usage history, as well as the used capacity, total capacity, and details of the space savings from deduplication and compression, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Volume Capacity Used (%)

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

Snapshot Capacity Used (%)

Displays the Snapshot reserve and Snapshot warning threshold as line graphs, and the capacity used by the Snapshot copies as an area graph, in percentage, on the vertical (y) axis. The Snapshot overflow is represented with different colors. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Snapshot Reserve legend, the Snapshot Reserve graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity

Displays the severity of the event.

Event

Displays the event name.

Triggered Time

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

Related Annotations pane

The Related Annotations pane enables you to view annotation details associated with the selected volume. The details include the annotation name and the annotation values that are applied to the volume. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view and navigate to the SVMs, aggregates, qtrees, LUNs, and Snapshot copies that are related to the volume:

Storage Virtual Machine

Displays the capacity and the health status of the SVM that contains the selected volume.

Aggregate

Displays the capacity and the health status of the aggregate that contains the selected volume. For FlexGroup volumes, the number of aggregates that comprise the FlexGroup is listed.

Volumes in the Aggregate

Displays the number and capacity of all the volumes that belong to the parent aggregate of the selected volume. The health status of the volumes is also displayed, based on the highest severity level. For example, if an aggregate contains ten volumes, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical. This component does not appear for FlexGroup volumes.

Qtrees

Displays the number of qtrees that the selected volume contains and the capacity of qtrees with quota that the selected volume contains. The capacity of the qtrees with quota is displayed in relation to the volume data capacity. The health status of the qtrees is also displayed, based on the highest severity level. For example, if a volume has ten qtrees, five with Warning status and the remaining five with Critical status, then the status displayed is Critical.

NFS 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

Performing diagnostic actions for volume offline conditions on page 33
After navigating to the Health/Volume details page of a volume reported to be offline, you can search for additional information helpful to diagnosing the volume offline condition.

Performing suggested remedial actions for a full volume on page 37
After receiving a "Volume Space Full" error event, you review the suggested remedial actions on the Event details page and decide to perform one of the suggested actions.

Health/Storage Virtual Machine details page

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

Note: You can monitor only data SVMs.

- *Command buttons*
- *Health tab*
- *Capacity tab*
- *Configuration tab*
- *LIFs tab*
- *Qtrees tab*
- *User and Group Quotas tab*
- *NFS Exports tab*
- *CIFS Shares tab*
- *SAN tab*
- *Data Policy tab*
- *Related Annotations pane*
- *Related Devices pane*
- *Related Groups pane*
- *Related Alerts pane*

Command buttons

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

Switch to Performance View

Enables you to navigate to the Performance/SVM details page.

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 Health/Storage Virtual Machines inventory 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 Health/Volumes inventory 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 Protection/Volume 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 or FlexGroup volume:

Capacity

The Capacity area displays details about the used and available capacity allocated from all volumes:

- Total Capacity
  Displays the total capacity (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 an indicated severity.
  You can move the pointer over the status to view more information about the capacity-related event or events generated for the volume.
  If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use the **View Details** button to view more information about the event.
  If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

  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.
- Aggregates
  For FlexVol volumes, displays the name of the aggregate that contains the volume. For FlexGroup volumes, displays the number of aggregates that are used in the FlexGroup.

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, FlexVol, or FlexVol/FlexGroup.
- Root Volume
  Displays the name of the root volume of the SVM.
- Allowed Protocols
  Displays the type of protocols that can be configured on the SVM. Also, indicates if a protocol is up (), down (), or is not configured ().

Data 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 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 command button

Opens the Edit Email Address dialog box, which displays the current email address of the selected user or user group. You can modify the email address. If the **Edit Email Address** field is blank, the default rule is used to generate an email address for the selected user or user group.

If more than one user has the same quota, the names of the users are displayed as comma-separated values. Also, the default rule is not used to generate the email address; therefore, you must provide the required email address for notifications to be sent.

Configure Email Rules command button

Enables you to create or modify rules to generate an email address for the user or user group quotas that are configured on the SVM. A notification is sent to the specified email address when there is a quota breach.

Status

Displays the current status of the quota. The status can be Critical (), Warning (), or Normal ().

You can move the pointer over the status icon to view more information about the event or events generated for the quota.

If the status of the quota is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the quota is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

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 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, FlexGroup 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 ONTAP version 8.2, 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, Restricted, or Mixed.

- Offline
  Read or write access to the volume is not allowed.

- Online
  Read and write access to the volume is allowed.
- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
- Mixed
  The constituents of a FlexGroup volume are not all in the same state.

Security Style

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

- UNIX (NFS clients)
  Files and directories in the volume have UNIX permissions.
- Unified
  Files and directories in the volume have a unified security style.
- NTFS (CIFS clients)
  Files and directories in the volume have Windows NTFS permissions.
- Mixed
  Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

UNIX Permission

Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.

Export Policy

Displays the rules that define the access permission for volumes that are exported. You can click the link to view details about the rules associated with the export policy such as the authentication protocols and the access permission.

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.

Note: Shares in folders are not displayed in the CIFS Shares tab.

View User Mapping command button

Launches the User Mapping dialog box.

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

Show ACL command button

Launches the Access Control dialog box for the share.

You can view user and permission details for the selected share.

Status

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

Share Name

Displays the name of the 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 ONTAP version 8.2, 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, Restricted, or Mixed.

- Offline
  Read or write access to the volume is not allowed.
- Online
  Read and write access to the volume is allowed.
- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
- Mixed
  The constituents of a FlexGroup volume are not all in the same state.

Security

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

- UNIX (NFS clients)
  Files and directories in the volume have UNIX permissions.
- NTFS (CIFS clients)
  Files and directories in the volume have Windows NTFS permissions.
- Mixed
  Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

Export Policy

Displays the name of the export policy applicable to the share. If an export policy is not specified for the SVM, the value is displayed as Not Enabled.

You can click the link to view details about the rules associated with the export policy, such as access protocols and permissions. The link is disabled if the export policy is disabled for the selected SVM.

NFS Equivalent

Specifies whether there is an NFS equivalent for the share.

SAN tab

Displays details about LUNs, initiator groups, and initiators for the selected SVM. By default, the LUNs view is displayed. You can view details about the initiator groups in the Initiator Groups tab and details about initiators in the Initiators tab.

LUNs tab

Displays details about the LUNs that belong to the selected SVM. You can view information such as the LUN name, LUN state (online or offline), the name of the file system (volume or qtree) that contains the LUN, the type of host operating system, the total data capacity and serial number of the LUN. You can also view information whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group.

You can also view the initiator groups and initiators that are mapped to the selected LUN.

Initiator Groups tab

Displays details about initiator groups. You can view details such as the name of the initiator group, the access state, the type of host operating system that is used by all the initiators in the group, and the supported protocol. When you click the link in the access state column, you can view the current access state of the initiator group.

Normal

The initiator group is connected to multiple access paths.

Single Path

The initiator group is connected to a single access path.

No Paths

There is no access path connected to the initiator group.

You can view whether initiator groups are mapped to all the 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 tab

Displays the name and type of the initiator and the total number of initiator groups mapped to this initiator for the selected SVM.

You can also view the LUNs and initiator groups that are mapped to the selected initiator group.

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. When there are FlexGroup volumes in the SVM, the count also includes FlexGroups; it does not include FlexGroup constituents.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected SVM.

Related Alerts pane

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

Health/Cluster details page

The Health/Cluster details page provides detailed information about a selected cluster, 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*
- *Health tab*
- *Capacity tab*
- *Configuration tab*
- *MetroCluster Connectivity tab*
- *MetroCluster Replication tab*
- *LIFs tab*
- *Nodes tab*
- *Disks tab*
- *Related Annotations pane*
- *Related Devices pane*
- *Related Groups pane*
- *Related Alerts pane*

Command buttons

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

Switch to Performance View

Enables you to navigate to the Performance/Cluster details page.



Enables you to add the selected cluster to the Favorites dashboard.

Actions

- Add Alert: Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.
- Rediscover: Initiates a manual refresh of the cluster, which enables Unified Manager to discover recent changes to the cluster.

  If Unified Manager is paired with OnCommand Workflow Automation, the rediscovery operation also reacquires cached data from WFA, if any.

  After the rediscovery operation is initiated, a link to the associated job details is displayed to enable tracking of the job status.
- Annotate: Enables you to annotate the selected cluster.

View Clusters

Enables you to navigate to the Health/Clusters inventory 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 the data capacity graph 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.

External Capacity Tier

Displays capacity details about the external capacity tier for FabricPool-enabled aggregates on the cluster. A FabricPool can be either licensed or unlicensed.

- Used
  Displays the space used by data in configured external capacity tiers.
- Data graph
  For an Amazon S3 FabricPool, the chart displays the total data capacity that has been licensed by this cluster and the amount being used by aggregates.
  For a StorageGRID FabricPool, the chart displays only the total capacity being used by aggregates.

Details

Displays detailed information about the used and available capacity.

- 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.
- Provisioned
  Displays the capacity (in MB, GB, and so on) that is provisioned for all the underlying volumes.
- Spares
  Displays the storable capacity (in MB, GB, and so on) available for storage in all the spare disks.
- External Capacity Tier
  Displays the space used by data in configured external capacity tiers. For an Amazon S3 FabricPool, the total data capacity that has been licensed by this cluster is also displayed.

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 ONTAP version that the cluster is running. If the nodes in the cluster are running different versions of ONTAP, then the earliest ONTAP version is displayed.
- Serial Number
  Displays the serial number of the cluster.
- Contact
  Displays details about the administrator whom you should contact in case of issues with the cluster.
- Location
  Displays the location of the cluster.

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 ONTAP versions that the nodes are running as well as the number of nodes running a particular version of ONTAP. For example, 9.0 (2), 8.3 (1) specifies that two nodes are running ONTAP 9.0, and one node is running ONTAP 8.3.

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.

External Capacity Tiers

Lists the names of the external capacity tiers to which this cluster is connected. It also lists the type (Amazon S3 or StorageGRID), and the states of the capacity tiers (Available or Unavailable).

MetroCluster Connectivity tab

Displays the issues and connectivity status of the cluster components in the MetroCluster configuration. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.

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.

Note: When managing an eight-node MetroCluster configuration, clicking the count link of the Disk Shelves component displays only the local shelves of the default HA pair. Also, there is no way to display the local shelves on the other HA pair.

You can move the pointer over the components to view the details and the connectivity status of the clusters in case of any issue and to view more information about the event or events generated for the issue.

If the status of the connectivity issue between components is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. The View Details button provides more information about the event.

If status of the connectivity issue between components is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

MetroCluster Replication tab

Displays the status of the data that is being replicated. You can use the MetroCluster Replication tab to ensure data protection by synchronously mirroring the data with the already peered clusters. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.

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 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 Health/Aggregates inventory 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 ONTAP.

Certain ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and spare management. ONTAP assigns an effective disk type for each disk type.

Spare Blocks Consumed %

Displays in percentage the spare blocks that are consumed in the SSD disk. This column is blank for disks other than SSD disks.

Rated Life Used %

Displays in percentage an estimate of the SSD life used, based on the actual SSD usage and the manufacturer's prediction of SSD life. A value greater than 99 indicates that the estimated endurance has been consumed, but may not indicate SSD failure. If the value is unknown, then the disk is omitted.

Firmware

Displays the firmware version of the disk.

RPM

Displays the revolutions per minute (RPM) of the disk. By default, this column is hidden.

Model

Displays the model number of the disk. By default, this column is hidden.

Vendor

Displays the name of the disk vendor. By default, this column is hidden.

Shelf ID

Displays the ID of the shelf where the disk is located.

Bay

Displays the ID of the bay where the disk is located.

Related Annotations pane

Enables you to view the annotation details associated with the selected cluster. The details include the annotation name and the annotation values that are applied to the cluster. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

Enables you to view device details that are associated with the selected cluster.

The details include properties of the device that is connected to the cluster such as the device type, size, count, and health status. You can click on the count link for further analysis on that particular device.

You can use MetroCluster Partner pane to obtain count and also details on the remote MetroCluster partner along with its associated cluster components such as nodes, aggregates, and SVMs. The MetroCluster Partner pane is displayed only for clusters in a MetroCluster configuration.

The Related Devices pane enables you to view and navigate to the nodes, SVMs, and aggregates that are related to the cluster:

MetroCluster Partner

Displays the health status of the MetroCluster partner. Using the count link, you can navigate further and obtain information about the health and capacity of the cluster components.

Nodes

Displays the number, capacity, and health status of the nodes that belong to the selected cluster. Capacity indicates the total usable capacity over available capacity.

Storage Virtual Machines

Displays the number of SVMs that belong to the selected cluster.

Aggregates

Displays the number, capacity, and the health status of the aggregates that belong to the selected cluster.

Related Groups pane

Enables you to view the list of groups that includes the selected cluster.

Related Alerts pane

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

Health/Aggregate details page

You can use the Health/Aggregate details page to view detailed information about the selected aggregate, such as the capacity, disk information, configuration details, and events generated. You can also view information about the related objects and related alerts for that aggregate.

- *Command buttons*
- *Capacity tab*
- *Disk Information tab*
- *Configuration tab*
- *History area*
- *Events list*
- *Related Devices pane*
- *Related Alerts pane*

Command buttons

Note: When monitoring a FabricPool-enabled aggregate, the committed and overcommitted values on this page are relevant only to the local, or performance tier, capacity. The amount of space available in the external capacity tier is not reflected in the overcommitted values. Similarly, the aggregate threshold values are relevant only to the local performance tier.

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

Switch to Performance View

Enables you to navigate to the Performance/Aggregate details page.



Enables you to add the selected aggregate to the Favorites dashboard.

Actions

- Add Alert
  Enables you to add an alert to the selected aggregate.

- Edit Thresholds
  Enables you to modify the threshold settings for the selected aggregate.

View Aggregates

Enables you to navigate to the Health/Aggregates inventory 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:

- 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.

External Capacity Tier

Displays capacity details about the external capacity tier for FabricPool-enabled aggregates. A FabricPool can be either licensed or unlicensed.

- Used
  Displays the space used by data in the external capacity tier.
- Unavailable
  Displays the space in the external capacity tier for an Amazon S3 FabricPool that cannot be used. This space may be shared with another FabricPool-enabled aggregate.
- Data graph
  For an Amazon S3 FabricPool, the chart displays the total data capacity that has been licensed by this cluster, the amount being used by this aggregate, and the unusable amount from other aggregates that are using the external capacity tier.
  For a StorageGRID FabricPool, the chart displays only the total capacity being used by this aggregate.

Details

Displays detailed information about capacity.

- Total Capacity
  Displays the total capacity in the aggregate.
- Data Capacity
  Displays the amount of space used by the aggregate (used capacity) and the amount of available space in the aggregate (free capacity).
- Snapshot Reserve
  Displays the used and free Snapshot capacity of the aggregate.
- Overcommitted Capacity
  Displays the aggregate overcommitment. Aggregate overcommitment enables you to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. When thin provisioning is in use, the total size of volumes in the aggregate can exceed the total capacity of the aggregate.

  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.
- External Capacity Tier
  For an Amazon S3 FabricPool, displays the total licensed capacity, the amount used by this aggregate, the amount used by other aggregates, and the free capacity for the external capacity tier. For a StorageGRID FabricPool, displays only the total capacity being used by this aggregate.
- 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.
- Other Details: Daily Growth Rate
  Displays the disk space used in the aggregate if the rate of change between the last two samples continues for 24 hours.
  For example, if an aggregate uses 10 GB of disk space at 2 pm and 12 GB at 6 pm, the daily growth rate (GB) for this aggregate is 2 GB.
- Volume Move
  Displays the number of volume move operations that are currently in progress:
  - Volumes Out
    Displays the number and capacity of the volumes that are being moved out of the aggregate.

You can click the link to view more details, such as the volume name, aggregate to which the volume is moved, status of the volume move operation, and the estimated end time.

- Volumes In

  Displays the number and remaining capacity of the volumes that are being moved into the aggregate.

  You can click the link to view more details, such as the volume name, aggregate from which the volume is moved, status of the volume move operation, and the estimated end time.

- Estimated used capacity after volume move

  Displays the estimated amount of used space (as a percentage, and in KB, MB, GB, and so on) in the aggregate after the volume move operations are complete.

Capacity Overview - Volumes

Displays graphs that provide information about the capacity of the volumes contained in the aggregate. The amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume is displayed. When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

You can select the graph you want to view from the drop-down lists. You can sort the data displayed in the graph to display details such as the used size, provisioned size, available capacity, fastest daily growth rate, and slowest growth rate. You can filter the data based on the storage virtual machines (SVMs) that contain the volumes in the aggregate. You can also view details for thinly provisioned volumes. You can view the details of specific points on the graph by positioning your cursor over the area of interest. By default, the graph displays the top 30 filtered volumes in the aggregate.

Disk Information tab

Displays detailed information about the disks in the selected aggregate, including the RAID type and size, and the type of disks used in the aggregate. The tab also graphically displays the RAID groups, and the types of disks used (such as SAS, ATA, FCAL, SSD, or VMDISK). You can view more information, such as the disk's bay, shelf, and rotational speed, by positioning your cursor over the parity disks and data disks.

Data

Graphically displays details about dedicated data disks, shared data disks, or both. When the data disks contain shared disks, graphical details of the shared disks are displayed. When the data disks contain dedicated disks and shared disks, graphical details of both the dedicated data disks and the shared data disks are displayed.

RAID Details

RAID details are displayed only for dedicated disks.

- Type
  Displays the RAID type (RAID0, RAID4, RAID-DP, or RAID-TEC).
- Group Size
  Displays the maximum number of disks allowed in the RAID group.
- Groups
  Displays the number of RAID groups in the aggregate.

Disks Used

- Effective Type

Displays the types of data disks (for example, ATA, SATA, FCAL, SSD, or VMDISK) in the aggregate.

- Data Disks
  Displays the number and capacity of the data disks that are assigned to an aggregate. Data disk details are not displayed when the aggregate contains only shared disks.
- Parity Disks
  Displays the number and capacity of the parity disks that are assigned to an aggregate. Parity disk details are not displayed when the aggregate contains only shared disks.
- Shared Disks
  Displays the number and capacity of the shared data disks that are assigned to an aggregate. Shared disk details are displayed only when the aggregate contains shared disks.

Spare Disks

Displays the disk effective type, number, and capacity of the spare data disks that are available for the node in the selected aggregate.

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.

External Capacity Tier

Displays the name of the external capacity tier, if you have configured a FabricPool-enabled aggregate, and shows the total licensed capacity for Amazon S3 FabricPools.

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.

External Capacity Tier

If this is a FabricPool-enabled aggregate, the details for the object store are displayed:

- Name
  Displays the name of the object store when it was created by ONTAP.
- Object Storage Provider
  Displays the name of the storage provider, for example, StorageGRID or Amazon S3.
- Object Store name (FQDN)
  Displays the FQDN of the object store.
- Access Key
  Displays the access key for the object store.
- Bucket Name
  Displays the bucket name of the object store.
- SSL
  Displays whether SSL encryption is enabled for the object store.

History area

The History area displays graphs that provide information about the capacity of the selected aggregate. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if the aggregate usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Aggregate Capacity Used (%)

Displays the used capacity in the aggregate and the trend in how aggregate capacity is used based on the usage history as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Capacity Used legend, the Capacity Used graph line is hidden.

Aggregate Capacity Used vs Total Capacity

Displays the trend in how aggregate capacity is used based on the usage history, as well as the used capacity and the total capacity, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Aggregate Capacity Used (%) vs Committed (%)

Displays the trend in how aggregate capacity is used based on the usage history, as well as the committed space as line graphs, as a percentage, on the vertical (y) axis. The time

period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Space Committed legend, the Space Committed graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity

Displays the severity of the event.

Event

Displays the event name.

Triggered Time

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp for when the event was generated is displayed.

Related Devices pane

The Related Devices pane enables you to view the cluster node, volumes, and disks that are related to the aggregate:

Node

Displays the capacity and the health status of the node that contains the aggregate. Capacity indicates the total usable capacity over available capacity.

Aggregates in the Node

Displays the number and capacity of all the aggregates in the cluster node that contains the selected aggregate. The health status of the aggregates is also displayed, based on the highest severity level. For example, if a cluster node contains ten aggregates, five of which display the Warning status and the remaining five of which display the Critical status, then the status displayed is Critical.

Volumes

Displays the number and capacity of FlexVol volumes and FlexGroup volumes in the aggregate; the number does not include FlexGroup constituents. The health status of the volumes is also displayed, based on the highest severity level.

Resource Pool

Displays the resource pools related to the aggregate.

Disks

Displays the number of disks in the selected aggregate.

Related Alerts pane

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

Related tasks

Performing suggested remedial actions for a full volume on page 37

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

Protection/Job details page

The Protection/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 Protection/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.

Definitions of user roles

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".

Integration Schema

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

Report Schema

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

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

Definitions of user types

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

Unified Manager user roles and capabilities

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

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

Function	Operator	Storage Administrator	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	•	•	•		
View threshold policies	•	•	•		
Manage storage service objects, such as SVM associations and resource pools		•	•		
Define alerts		•	•		
Manage storage management options		•	•		
Manage storage management policies		•	•		
Manage users			•		
Manage administrative options			•		
Define threshold policies			•		
Manage database access			•		
Manage integration with WFA and provide access to the database views				•	
Provide read-only access to reporting and other database views					•
Schedule and save reports	•	•	•		
Import and delete imported reports			•		

Related reference

Definitions of user types on page 163

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.

Definitions of user roles on page 162

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.

Supported Unified Manager CLI commands

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

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

CLI command	Description	Output
`um run cmd [ -t <timeout> ] <cluster> <command>`	The simplest way to run a command on one or more hosts. Mainly used for alert scripting to get or perform an operation on ONTAP. The optional timeout argument sets a maximum time limit (in seconds) for the command to complete on the client. The default is 0 (wait forever).	As received from ONTAP.
`um run query <sql command>`	Executes an SQL query. Only queries that read from the database are allowed. Any update, insert, or delete operations are not supported.	Results are displayed in a tabular form. If an empty set is returned, or if there is any syntax error or bad request, it displays the appropriate error message.
`um datasource add -u <username> -P <password> [ -t <protocol> ] [ -p <port> ] <hostname-or-ip>`	Adds a datasource to the list of managed storage systems. A datasource describes how connections to storage systems are made. The options -u (username) and -P (password) must be specified when adding a datasource. The option -t (protocol) specifies the protocol used to communicate with the cluster (http or https). If the protocol is not specified, then both protocols will be attempted The option -p (port) specifies the port used to communicate with the cluster. If the port is not specified, then the default value of the appropriate protocol will be attempted. This command can be executed only by the storage admin.	Prompts for the user accept the certificate and prints the corresponding message.

CLI command	Description	Output
`um datasource list [ <datasource-id>]`	Displays the datasources for managed storage systems.	Displays the following values in tabular format: `ID Address Port, Protocol Acquisition Status, Analysis Status, Communication status, Acquisition Message, and Analysis Message` .
`um datasource modify [ -h <hostname-or-ip> ] [ -u <username> ] [ -P <password> ] [ -t <protocol> ] [ -p <port> ] <datasource-id>`	Modifies one or more datasource options. Can be executed only by the storage admin.	Displays the corresponding message.
`um datasource remove <datasource-id>`	Removes the datasource from Unified Manager.	Displays the corresponding message.
`um option list [ <option> .. ]`	Lists options.	Displays the following values in tabular format: `Name, Value, Default Value, and Requires Restart.`
`um option set <option-name>=<option-value> [ <option-name>=<option-value> ... ]`	Sets one or more options. The command can be executed only by the storage admin.	Displays the corresponding message.
`um version`	Displays the Unified Manager software version .	`Version ("7.0")`
`um lun list [-q] [ -ObjectType <object-id>]`	Lists the LUNs after filtering on the specified object. -q is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: **`um lun list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the LUNs within the cluster with ID 1.	Displays the following values in tabular format: `ID and LUN path` .

CLI command	Description	Output
`um svm list [-q] [ -ObjectType <object-id>]`	Lists the SVMs after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: **`um svm list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the SVMs within the cluster with ID 1.	Displays the following values in tabular format: `Name and Cluster ID` .
`um qtree list [-q] [ -ObjectType <object-id>]`	Lists the qtrees after filtering on the specified object. -q is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: **`um qtree list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the qtrees within the cluster with ID 1.	Displays the following values in tabular format: `Qtree ID and Qtree Name` .
`um disk list [-q] [-ObjectType <object-id>]`	Lists the disks after filtering on the specified object. ObjectType can be disk, aggr, node, cluster. For example: **`um disk list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the disks within the cluster with ID 1.	Displays the following values in tabular format `ObjectType and object-id.`
`um cluster list [-q] [-ObjectType <object-id>]`	Lists the clusters after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, lun, qtree, volume, quota, svm. For example: **`um cluster list -aggr 1`** In this example, "-aggr" is the objectType and "1" is the objectId. The command lists the cluster to which the aggregate with ID 1 belongs.	Displays the following values in tabular format: `Name, Full Name, Serial Number, Datasource Id, Last Refresh Time, and Resource Key.`
`um cluster node list [-q] [-ObjectType <object-id>]`	Lists the cluster nodes after filtering on the specified object. ObjectType can be disk, aggr, node, cluster. For example: **`um cluster node list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the nodes within the cluster with ID 1.	Displays the following values in tabular format `Name and Cluster ID.`

CLI command	Description	Output
`um volume list [-q] [-ObjectType <object-id>]`	Lists the volumes after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, svm, aggregate. For example: **`um volume list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the volumes within the cluster with ID 1.	Displays the following values in tabular format `Volume ID and Volume Name.`
`um quota user list [-q] [-ObjectType <object-id>]`	Lists the quota users after filtering on the specified object. ObjectType can be qtree, cluster, volume, quota, svm. For example: **`um quota user list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the quota users within the cluster with ID 1.	Displays the following values in tabular format `ID, Name, SID and Email.`
`um aggr list [-q] [-ObjectType <object-id>]`	Lists the aggregates after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, volume. For example: **`um aggr list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the aggregates within the cluster with ID 1.	Displays the following values in tabular format `Aggr ID, and Aggr Name.`
`um event ack <event-ids>`	Acknowledges one or more events.	Displays the corresponding message.
`um event resolve <event-ids>`	Resolves one or more events.	Displays the corresponding message.
`um event assign -u <username> <event-id>`	Assigns an event to a user.	Displays the corresponding message.
`um event list [ -s <source> ] [ -S <event-state-filter-list>.. ] [ <event-id> .. ]`	Lists the events generated by the system or user. Filters events based on source, state, and IDs.	Displays the following values in tabular format `Source, Source type, Name, Severity, State, User and Timestamp.`
`um cli login -u <username> [-p <password>`	Logs in to the CLI. The session expires after three hours from the time of login, after which the user must login again.	Displays the corresponding message.
`um cli logout`	Logs out of the CLI.	Displays the corresponding message.
`um backup restore -f <path of .7z file>`	Enables users to perform restore operation using a .7z file.	Displays the corresponding message.

CLI command	Description	Output
`um help`	Displays all first level subcommands.	Displays all first level subcommands.

Using the maintenance console

You can use the maintenance console to configure network settings, to configure and manage the system on which Unified Manager is installed, and to perform other maintenance tasks that help you prevent and troubleshoot possible issues.

Related concepts

What functionality the maintenance console provides on page 170
The Unified Manager maintenance console enables you to maintain the settings on your Unified Manager system and to make any necessary changes to prevent issues from occurring.

Diagnostic user capabilities on page 171
The purpose of diagnostic access is to enable technical support to assist you in troubleshooting, and you should only use it when directed by technical support.

Related tasks

Sending a Unified Manager support bundle to technical support on page 105
This workflow shows you how to generate, retrieve, and send a support bundle to technical support using the Unified Manager maintenance console. You should send a support bundle when the issue that you have requires more detailed diagnosis and troubleshooting than an AutoSupport message provides.

What functionality the maintenance console provides

The Unified Manager maintenance console enables you to maintain the settings on your Unified Manager system and to make any necessary changes to prevent issues from occurring.

Depending on the operating system on which you have installed Unified Manager, the maintenance console provides the following functions:

- Troubleshoot any issues with your virtual appliance, especially if the Unified Manager web interface is not available
- Upgrade to newer versions of Unified Manager
- Generate support bundles to send to technical support
- Configure network settings
- Change the maintenance user password
- Connect to an external data provider to send performance statistics
- Change the performance data collection internal
- Restores the Unified Manager database and configuration settings from a previously backed up version.
- Import performance data from ONTAP clusters that have been managed by OnCommand Performance Manager 7.1.

Related tasks

Using the maintenance console on page 170
You can use the maintenance console to configure network settings, to configure and manage the system on which Unified Manager is installed, and to perform other maintenance tasks that help you prevent and troubleshoot possible issues.

What the maintenance user does

The maintenance user is created during the installation of Unified Manager. For Red Hat installations the maintenance user name is the "umadmin" user. The maintenance user has the

OnCommand administrator role in the web UI, and that user can create subsequent users and assign them roles.

The maintenance user, or umadmin user, can also access the Unified Manager maintenance console.

Related tasks

Using the maintenance console on page 170

You can use the maintenance console to configure network settings, to configure and manage the system on which Unified Manager is installed, and to perform other maintenance tasks that help you prevent and troubleshoot possible issues.

Diagnostic user capabilities

The purpose of diagnostic access is to enable technical support to assist you in troubleshooting, and you should only use it when directed by technical support.

The diagnostic user can execute OS-level commands when directed by technical support, for troubleshooting purposes.

Related tasks

Using the maintenance console on page 170

You can use the maintenance console to configure network settings, to configure and manage the system on which Unified Manager is installed, and to perform other maintenance tasks that help you prevent and troubleshoot possible issues.

Accessing the maintenance console

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

Before you begin

You must have installed and configured Unified Manager.

About this task

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

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

Step

Follow these steps to access the maintenance console:

On this operating system...	Follow these steps...
VMware	**a.** Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance. **b.** Log in to the maintenance console using your maintenance user name and password.
Red Hat	**a.** Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system. **b.** Log in to the system with the maintenance user (umadmin) name and password. **c.** Enter the command **`maintenance_console`** and press Enter.

On this operating system...	Follow these steps...
Windows	**a.** Log in to the Unified Manager system with administrator credentials. **b.** Launch PowerShell as a Windows administrator. **c.** Enter the command **`maintenance_console`** and press Enter. **Note:** On Microsoft Windows Server 2012 if you receive an execution policy error, enter the following command and try step c again: **`PowerShell.exe -ExecutionPolicy RemoteSigned`**

The Unified Manager maintenance console menu is displayed.

Related tasks

Using the maintenance console on page 170

You can use the maintenance console to configure network settings, to configure and manage the system on which Unified Manager is installed, and to perform other maintenance tasks that help you prevent and troubleshoot possible issues.

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

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.

Maintenance console menus

The maintenance console consists of different menus that enable you to maintain and manage special features and configuration settings of the Unified Manager server.

Depending on the operating system on which you have installed Unified Manager, the maintenance console consists of the following menus:

- Upgrade Unified Manager (VMware only)
- Network Configuration (VMware only)
- System Configuration (VMware only)
- Support/ Diagnostics
- Reset Server Certificate
- External Data Provider
- Performance Polling Interval Configuration

Network Configuration menu

The Network Configuration menu enables you to manage the network settings. You should use this menu when the OnCommand Unified Manager user interface is not available.

Note: This menu is not available if Unified Manager is installed on Red Hat Enterprise Linux or on Microsoft Windows.

The following menu choices are available.

Display IP Address Settings

Displays the current network settings for the virtual appliance, including the IP address, network, broadcast address, netmask, gateway, and DNS servers.

Change IP Address Settings

Enables you to change any of the network settings for the virtual appliance, including the IP address, netmask, gateway, or DNS servers. If you switch your network settings from DHCP to static networking using the maintenance console, you cannot edit the host name. You must select **Commit Changes** for the changes to take place.

Display Domain Name Search Settings

Displays the domain name search list used for resolving host names.

Change Domain Name Search Settings

Enables you to change the domain names for which you want to search when resolving host names. You must select **Commit Changes** for the changes to take place.

Display Static Routes

Displays the current static network routes.

Change Static Routes

Enables you to add or delete static network routes. You must select **Commit Changes** for the changes to take place.

Add Route

Enables you to add a static route.

Delete Route

Enables you to delete a static route.

Back

Takes you back to the **Main Menu**.

Exit

Exits the maintenance console.

Disable Network Interface

Disables any available network interfaces. If only one network interface is available, you cannot disable it. You must select **Commit Changes** for the changes to take place.

Enable Network Interface

Enables available network interfaces. You must select **Commit Changes** for the changes to take place.

Commit Changes

Applies any changes made to the network settings for the virtual appliance. You must select this option to enact any changes made, or the changes do not occur.

Ping a Host

Pings a target host to confirm IP address changes or DNS configurations.

Restore to Default Settings

Resets all settings to the factory default. You must select **Commit Changes** for the changes to take place.

Back

Takes you back to the **Main Menu**.

Exit

Exits the maintenance console.

System Configuration menu

The System Configuration menu enables you to manage your virtual appliance by providing various options, such as viewing the server status, and rebooting and shutting down the virtual machine.

Note: The System Configuration menu is not available if Unified Manager is installed on Red Hat Enterprise Linux or Microsoft Windows.

The following menu choices are available:

Display Server Status

Displays the current server status. Status options include Running and Not Running.

If the server is not running, you might need to contact technical support.

Reboot Virtual Machine

Reboots the virtual machine, stopping all services. After rebooting, the virtual machine and services restart.

Shut Down Virtual Machine

Shuts down the virtual machine, stopping all services.

You can select this option only from the virtual machine console.

Change <logged in user> User Password

Changes the password of the user that is currently logged in, which can only be the maintenance user.

Increase Data Disk Size

Increases the size of the data disk (disk 3) in the virtual machine.

Increase Swap Disk Size

Increases the size of the swap disk (disk 2) in the virtual machine.

Change Time Zone

Changes the time zone to your location.

Change NTP Server

Changes the NTP Server settings, such as IP address or fully qualified domain name (FQDN).

Restore from an OCUM Backup

Restores the Unified Manager database and configuration settings from a previously backed up version.

Reset Server Certificate

Resets the server security certificate.

Change hostname

Changes the name of the host on which the virtual appliance is installed.

emu

Enables the migration tool user to migrate monitoring health data from Unified Manager 5.x to the current version of Unified Manager.

This option is hidden; you must type the entry.

Disable Migration Tool user

Disables the migration tool user - only displayed if the migration tool user is enabled.

Back

Exits the System Configuration menu and returns to the Main Menu.

Exit

Exits the maintenance console menu.

Support and Diagnostics menu

The Support and Diagnostics menu enables you to generate a support bundle.

The following menu option is available:

Generate Support Bundle

Enables you to create a 7-Zip file containing full diagnostic information in the diagnostic user's home directory. The file includes information generated by an AutoSupport message, the contents of the Unified Manager database, detailed data about the Unified Manager server internals, and verbose-level logs not normally included in AutoSupport messages.

Additional menu options

The following menu options enable you to perform various administrative tasks on the Unified Manager server.

The following menu choices are available:

Reset Server Certificate

Regenerates the HTTPS server certificate.

You can regenerate the server certificate in the Unified Manager GUI by clicking  > **HTTPS Certificate** > **Regenerate HTTPS Certificate**.

Disable SAML authentication

Disables SAML authentication so that the identity provider (IdP) no longer provides sign-on authentication for users accessing the Unified Manager GUI. This console option is typically used when an issue with the IdP server or SAML configuration blocks users from accessing the Unified Manager GUI.

External Data Provider

Provides options for connecting Unified Manager to an external data provider. After you establish the connection, performance data is sent to an external server so that storage performance experts can chart the performance metrics using third-party software. The following options are displayed:

- **Display Server Configuration**—Displays the current connection and configuration settings for an external data provider.
- **Add / Modify Server Connection**—Enables you to enter new connection settings for an external data provider, or change existing settings.

- **Modify Server Configuration**—Enables you to enter new configuration settings for an external data provider, or change existing settings.
- **Delete Server Connection**—Deletes the connection to an external data provider.
  After the connection is deleted, Unified Manager loses its connection to the external server.

Performance Polling Interval Configuration

Provides an option for configuring how frequently Unified Manager collects performance statistical data from clusters. The default collection interval is five minutes.

You can change this interval to ten or fifteen minutes if you find that collections from large clusters are not completing on time.

Exit

Exits the maintenance console menu.

Changing the maintenance user password on Windows

You can change the Unified Manager maintenance user password when required.

Steps

1. From the Unified Manager web UI login page, click **Forgot Password**.
   A page is displayed that prompts for the name of the user whose password you want to reset.
2. Enter the user name and click **Submit**.
   An email with a link to reset the password is sent to the email address that is defined for that user name.
3. Click the **reset password link** in the email and define the new password.
4. Return to the web UI and log in to Unified Manager using the new password.

After you finish

If Unified Manager is installed in a Microsoft Cluster Server (MSCS) environment, then you must change the maintenance user password on the second node of the MSCS setup. The maintenance user password for both nodes must be same.

Changing the umadmin password on Red Hat Enterprise Linux

For security reasons, you must change the default password for the Unified Manager umadmin user immediately after completing the installation process. If necessary, you can change the password again anytime later. On Red Hat Enterprise Linux, you can change the umadmin password with a Linux command.

Before you begin

- Unified Manager must be installed on a Red Hat Enterprise Linux system.
- You must have the root user credentials for the Red Hat Enterprise Linux system on which Unified Manager is installed.

Steps

1. Log in as the root user to the Red Hat Enterprise Linux system on which Unified Manager is running.
2. Change the umadmin password:

   ```
   passwd umadmin
   ```

   The system prompts you to enter a new password for the umadmin user.

Note: If Unified Manager is installed in a Veritas Cluster Server (VCS) environment, you must change the umadmin password on both nodes of the VCS setup. The umadmin password for both nodes must be the same.

Adding network interfaces

You can add new network interfaces if you need to separate network traffic.

Before you begin

You must have added the network interface to the virtual appliance using vSphere.

The virtual appliance must be powered on.

About this task

Note: You cannot perform this operation if Unified Manager is installed on Red Hat Enterprise Linux or on Microsoft Windows.

Steps

1. In the vSphere console **Main Menu**, select **System Configuration** > **Reboot Operating System**.

   After rebooting, the maintenance console can detect the newly added network interface.
2. Access the maintenance console.
3. Select **Network Configuration** > **Enable Network Interface**.
4. Select the new network interface and press **Enter**.
   Select **eth1** and press **Enter**.
5. Type **y** to enable the network interface.
6. Enter the network settings.

   You are prompted to enter the network settings if using a static interface, or if DHCP is not detected.

   After entering the network settings, you automatically return to the **Network Configuration** menu.
7. Select **Commit Changes**.

   You must commit the changes to add the network interface.

Related tasks

Accessing the maintenance console using the vSphere VM console

Adding disk space to the Unified Manager database directory

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

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

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

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

Adding space to the data directory of the Red Hat Enterprise Linux host

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

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

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

Steps

1. Log in as root user to the Red Hat Enterprise Linux machine on which you want to add disk space.
2. Stop the Unified Manager service and the associated MySQL software in the order shown:

   ```
   service ocieau stop
   service ocie stop
   service mysqld stop
   ```
3. Create a temporary backup folder (for example, `/backup-data`) with sufficient disk space to contain the data in the current `/opt/netapp/data` directory.
4. Copy the content and privilege configuration of the existing `/opt/netapp/data` directory to the backup data directory:

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

      ```
      se_type=`ls -Z /opt/netapp/data | awk '{print $4}' | awk -F: '{print $3}' | head -1`
      ```
      The system returns a confirmation similar to the following:

      ```
      echo $se_type
      mysqld_db_t
      ```
   b. Run the `chcon` command to set the SE Linux type for the backup directory:

      ```
      chcon -R --type=mysqld_db_t /backup-data
      ```
6. Remove the contents of the `/opt/netapp/data` directory:

   ```
   cd /opt/netapp/data
   rm -rf *
   ```
7. Expand the size of the `/opt/netapp/data` directory to a minimum of 750 GB through LVM commands or by adding extra disks.

   Important: Mounting the `/opt/netapp/data` directory on an NFS export or CIFS share is not supported.
8. Confirm that the `/opt/netapp/data` directory owner (mysql) and group (root) are unchanged:

   ```
   ls -ltr / | grep opt/netapp/data
   ```
   The system returns a confirmation similar to the following:

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

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

   `touch /opt/netapp/data/abc`

   `ls -Z /opt/netapp/data/abc`
   The system returns a confirmation similar to the following:

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

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

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

11. Start the MySQL service:

    `service mysqld start`

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

    `service ocie start`

    `service ocieau start`

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

    `rm -rf /backup-data`

Adding space to the data disk of the VMware virtual machine

If you need to increase the amount of space on the data disk for the Unified Manager database, you can add capacity after installation by increasing disk space on `disk 3`.

Before you begin

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

About this task

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

Steps

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

Adding space to the logical drive of the Microsoft Windows server

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

Before you begin

You must have Windows administrator privileges.

About this task

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

Steps

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

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

Copyright and trademark

Copyright

Copyright © 2020 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

http://www.netapp.com/us/legal/netapptmlist.aspx