OnCommand® Unified Manager 6.1

Online Help

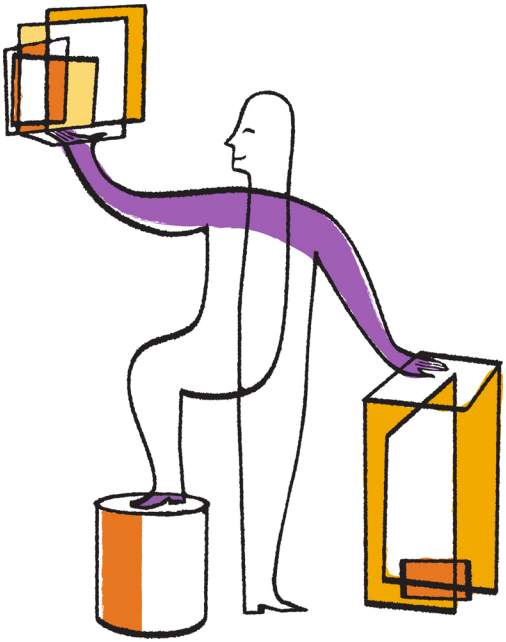

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-08261_A0
February 2014

Contents

Introduction to OnCommand Unified Manager 13
OnCommand Unified Manager features 13
What you can do with Unified Manager 14
Understanding the workspace 15
Typical window layouts 15
Window layout customization 16
Using the Unified Manager Help 17
Bookmarking your favorite topics 18
Exporting data to CSV files 18
What the Unified Manager server does 20
How the discovery process works 20
Searching for storage objects 21
How timestamps work in Unified Manager 21
Getting started 22
Configuring your environment after deployment 22
Changing the Unified Manager host name 23
Configuring Unified Manager to send alert notifications 24
Customizing your environment 25
Enabling periodic AutoSupport 25
Editing the network settings 25
Editing the global setup options 26
Working with HTTPS security certificates 27
Description of Setup options dialog boxes 28
Setup Options dialog box 28
Configure Network Settings dialog box 31
AutoSupport Setup Options dialog box 32
Monitoring your system from the dashboard 33
Understanding the dashboard 33
Description of the Dashboard window 35
Cluster availability details pane 35
Quick Takes area 35
Unresolved Incidents and Risks area 36

Monitoring performance 38
Enabling performance monitoring 38
Creating a user with Event Publisher role privileges 38
Configuring a connection between a Performance Manager server and the Unified Manager server 39
Deleting a connection between a Performance Manager server and the Unified Manager server 40
Viewing details of performance incidents listed in the Dashboard 41
Viewing all performance incidents 42
Viewing details of performance incidents listed on the Events page 42
Understanding more about performance monitoring 43
What performance incidents are 43
Purpose of a connection between Performance Manager and Unified Manager 44
Connections between multiple Performance Manager servers and Unified Manager 44
Display of performance incidents in Unified Manager 44
Workloads, cluster components, and performance 45
How performance incidents are discovered and analyzed 47
Managing events and alerts 48
What events are 48
Configuring event settings 48
Configuring notification settings 49
Viewing event details 50
Disabling or enabling events 51
Assigning events 52
Viewing unassigned events 53
Adding and reviewing notes about an event 53
Acknowledging and resolving events 54
Understanding more about events 55
Event state definitions 55
Description of event severity types 55
Description of event impact levels 56
Description of event impact areas 56
How object status is computed 57
List of events and severity types 57

What alerts are ... 77
Adding an alert ... 77
Guidelines for adding alerts ... 79
Testing an alert ... 80
Viewing alerts ... 80
Editing an alert ... 81
Deleting alerts ... 81
Description of Event windows and dialog boxes ... 82
Event Retention Settings ... 82
Notification Setup Options dialog box ... 83
Events page ... 85
Event details page ... 87
Manage Events page ... 90
Disable Events dialog box ... 91
Description of Alert windows and dialog boxes ... 92
Alerts page ... 92
Add Alert dialog box ... 93
Edit Alert dialog box ... 95
Managing thresholds ... 98
What storage capacity thresholds are ... 98
Configuring global aggregate threshold values ... 98
Configuring global volume threshold values ... 99
Editing unmanaged relationship lag threshold settings ... 99
Editing global threshold settings ... 100
Editing aggregate threshold settings ... 101
Editing volume threshold settings ... 102
Description of thresholds dialog boxes ... 102
Aggregate Thresholds Setup Options dialog box ... 102
Volume Thresholds Setup Options dialog box ... 105
Managing quotas ... 109
What quota limits are ... 109
Viewing user and user group quotas ... 109
Creating rules to generate email addresses ... 110
Creating an email notification format for user and user group quotas ... 111
Editing user and group quota email addresses ... 111
Understanding more about quotas ... 112

Overview of the quota process 113
About quotas 113
Why you use quotas 113
Description of quotas dialog boxes 114
Email Notification Format dialog box 114
Rules to Generate User and Group Quota Email Address dialog box 115
Managing and monitoring clusters and cluster objects 118
Understanding cluster monitoring 118
Understanding node root volumes 119
Understanding capacity events and thresholds for node root aggregates 119
Adding clusters 120
Removing a cluster in case of discovery failure 121
Restarting cluster discovery 122
Editing clusters 122
Removing clusters 123
Rediscovering updated clusters manually 123
Viewing the cluster list and details 124
Viewing the node list and details 125
Viewing the SVM list and details 126
Viewing the aggregate list and details 127
Viewing the volume list and details 128
Viewing the CIFS shares 129
Viewing the data sources 129
Viewing the list of Snapshot copies 130
Deleting Snapshot copies 130
Calculating reclaimable space for Snapshot copies 131
Understanding clusters and cluster objects 132
What a cluster is 132
What SVMs are 132
Why you use SVMs 134
What logical storage is 135
How volumes work 135
What LIFs are 136
Types of network ports 137
What the Snapshot copy reserve is 138
What an HA pair is 138

How HA pairs relate to the cluster 138
When takeovers occur 141
Connections and components of an HA pair 142
Understanding takeover and giveback 142
About CIFS and SMB 143
CIFS license 143
Supported SMB clients and domain controllers 143
What a Snapshot copy is 144
Description of cluster object windows and dialog boxes 144
Clusters page 144
Cluster details page 146
Add Cluster dialog box 155
Edit Cluster dialog box 156
Nodes page 156
Storage Virtual Machines page 158
Storage Virtual Machine details page 160
Aggregates page 175
Aggregate details page 178
Volumes page 184
Volume details page 190
Export Policy Rules dialog box 204
Data Sources page 205
Snapshot Copies on a Volume dialog box 206
Managing annotations for storage objects 208
What annotations are 208
Creating rules to annotate storage objects 208
Viewing annotations 209
Removing annotations for storage objects 209
Understanding more about annotations 210
Description of annotation types 210
Description of Annotations windows and dialog boxes 210
Manage Annotations page 210
Configure Workflow Annotation Rules dialog box 212
Managing and monitoring protection relationships 215
What resource pools are 215
Viewing the protection inventory 215

Creating a SnapVault protection relationship from the Volumes page 216
Creating a SnapVault protection relationship from the Volume details page 217
Creating a SnapMirror protection relationship from the Volume details page 218
Creating a SnapMirror protection relationship from the Volumes page 219
Configuring destination efficiency settings 221
Creating SnapMirror and SnapVault schedules 221
Editing protection relationships from the Volumes page 222
Editing protection relationships from the Volume details page 223
Creating SnapMirror policies 224
Creating SnapVault policies 224
Aborting an active data protection transfer from the Volume details page 225
Aborting an active data protection transfer from the Volumes page 226
Quiescing protection relationships from the Volume details page 227
Breaking a SnapMirror relationship from the Volume details page 228
Removing a protection relationship from the Volume details page 229
Resuming scheduled transfers on a quiesced relationship from the Volumes page 229
Resuming scheduled transfers on a quiesced relationship from the Volume details page 230
Initializing or updating protection relationships from the Volume details page 231
Initializing or updating protection relationships from the Volumes page 232
Resynchronizing protection relationships from the Volume details page 233
Resynchronizing protection relationships from the Volumes page 234
Reversing protection relationships from the Volume details page 235
Reversing protection relationships from the Volumes page 236
Restoring data using the Volumes page 238
Restoring data using the Volume details page 239
Creating resource pools 240
Editing resource pools 241
Viewing resource pools inventory 241
Adding resource pool members 242
Removing aggregates from resource pools 243
Deleting resource pools 243
Understanding SVM associations 244
Storage Virtual Machine (SVM) and resource pool requirements to support storage services 245

Creating Storage Virtual Machine (SVM) associations 245
Viewing SVM associations 247
Deleting SVM associations 247
What jobs are 248
Monitoring jobs 248
Viewing job details 249
Aborting jobs 249
Retrying a failed protection job 250
Description of Protection relationships window and dialog boxes 250
Resource Pools page 251
Create Resource Pool dialog box 253
Edit Resource Pool dialog box 254
Aggregates dialog box 255
SVM Associations page 257
Create Storage Virtual Machine Associations wizard 258
Jobs page 259
Advanced Secondary Settings dialog box 261
Advanced Destination Settings dialog box 262
Restore dialog box 263
Browse Directories dialog box 265
Job details page 266
Relationships Thresholds Setup Options dialog box 267
Edit Aggregate Thresholds dialog box 268
Edit Volume Thresholds dialog box 270
Configure Protection dialog box 272
Create New Schedule dialog box 278
Create SnapMirror Policy dialog box 279
Create SnapVault Policy dialog box 280
Edit Relationship dialog box 281
Initialize/Update dialog box 282
Resynchronize dialog box 284
Select Source Snapshot Copy dialog box 285
Reverse Resync dialog box 288

Executing protection workflows using OnCommand Workflow Automation 290

Pairing OnCommand Workflow Automation with Unified Manager 290

Removing OnCommand Workflow Automation setup from Unified Manager 291
What happens when OnCommand Workflow Automation is reinstalled or upgraded 291
Description of OnCommand Workflow Automation setup windows and dialog boxes 292
Set Up OnCommand Workflow Automation dialog box 292
Managing and monitoring Infinite Volumes 294
Viewing the details of SVMs with Infinite Volume 294
Viewing the constituents of an Infinite Volume 294
Editing the Infinite Volume threshold settings 295
Editing the threshold settings of storage classes 296
Understanding Infinite Volumes 296
What an Infinite Volume is 297
Maximum number of files an Infinite Volume can store 297
What a storage class is 297
What a namespace constituent is 298
What data constituents are 298
What a namespace mirror constituent is 298
Creating rules 298
Creating rules using templates 299
Creating custom rules 300
Viewing rules 301
Editing template-based rules 302
Editing custom rules 303
Deleting rules 304
Previewing changes to your data policy 304
Exporting a data policy configuration 305
Importing a data policy configuration 306
Understanding rules and data policy 307
What rules and data policies are 307
What the default rule is 308
How a data policy filters data written to an Infinite Volume 308
What a rule template is 309
What conditions and condition sets are 309
Description of Infinite Volume windows and dialog boxes 310
Create Rule dialog box 310

Edit Rule dialog box ... 311
Edit Rule dialog box (Advanced edit) ... 312
Edit SVM with Infinite Volume Thresholds dialog box ... 313
Edit Storage Class Thresholds dialog box ... 314
Managing user access ... 316
Adding a user ... 316
Editing the user settings ... 317
Testing a remote user or group ... 317
Viewing users ... 318
Deleting users or groups ... 318
Changing the local user password ... 319
What the maintenance user does ... 320
What RBAC is ... 320
What role-based access control does ... 320
Definitions of user types ... 321
Definitions of user roles in Unified Manager ... 321
Unified Manager roles and capabilities ... 322
Description of user access windows and dialog boxes ... 323
Manage Users page ... 323
Add User dialog box ... 324
Edit User dialog box ... 325
Managing authentication ... 327
Enabling remote authentication ... 327
Setting up authentication services ... 328
Adding authentication servers ... 329
Testing the configuration of authentication servers ... 330
Editing authentication servers ... 331
Deleting authentication servers ... 331
Authentication with Active Directory or OpenLDAP ... 332
Description of authentication windows and dialog boxes ... 332
Authentication Setup Options dialog box ... 332
Managing security certificates ... 337
Viewing the HTTPS security certificate ... 337
Generating an HTTPS security certificate ... 338
Downloading an HTTPS certificate signing request ... 340
Installing an HTTPS security certificate ... 341

Description of security certificates windows and dialog boxes 342
HTTPS Setup Options dialog box 342
Troubleshooting 344
Sending an on-demand AutoSupport message 344
Unknown authentication error 344
User not found 345
Icons are misaligned in Internet Explorer 345
LDAP server slow to respond 346
Issue with adding LDAP using Other authentication services 346
Troubleshooting access to CIFS shares 347
Identifying inaccessible volumes and qtrees 348
Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring 348
Glossary 350
Copyright information 365
Trademark information 366
How to send your comments 367
Index 368

Introduction to OnCommand Unified Manager

OnCommand Unified Manager is part of the OnCommand family of management software. Unified Manager provides IT efficiency savings by unifying and automating storage operations, provisioning, and protection for both physical resources such as resource pools and virtual resources such as Storage Virtual Machines (SVMs).

If any issues occur in the cluster objects that are monitored, Unified Manager notifies you through events about such issues so that you can take the required corrective action.

OnCommand Unified Manager features

Unified Manager is built on a server infrastructure that delivers scalability, supportability, and enhanced monitoring and notification capabilities. Unified Manager supports monitoring of clustered Data ONTAP 8.1.1, 8.1.2, 8.1.3, 8.1.4, 8.2, and 8.2.1. However, for Data ONTAP 8.2 and earlier, Data Protection monitoring is not supported.

Unified Manager includes the following features:

- Discovery, monitoring, and notifications for clustered Data ONTAP:
  - Physical objects: cluster nodes, disks, disk shelves, SFO pairs, ports, and Flash Cache
  - Logical objects: clusters, Storage Virtual Machines (SVMs, formerly known as Vservers), aggregates, volumes, LUNs, qtrees, LIFs, Snapshot copies, junction paths, NFS exports, CIFS shares, user and group quotas, Initiator Groups, Initiator
  - Protocols: CIFS, NFS, FC, iSCSI, and FCoE
  - Storage efficiency: Flash Pool aggregates, deduplication, and compression
  - Protection: SnapMirror relationships and SnapVault relationships
- Enhanced performance monitoring
- Viewing the cluster discovery and monitoring status
- Enhanced UI and comprehensive cluster visualization
- Enhanced alerts, events, and threshold infrastructure
- LDAP and local user support
- RBAC (only for a predefined set of roles)
- AutoSupport and support bundle
- Enhanced logging
- Enhanced dashboard to show capacity, availability, protection, and performance health of the environment
- Volume move interoperability, volume move history, and junction path change history
- Remedial suggestions for taking actions on the Volume Space Nearly Full event and the Volume Space Full event

- Impacted resources are displayed for the Volume Offline event, the Volume Restricted event, and the Thin-Provisioned Volume Space At Risk event.
- Support for SVMs with Infinite Volume:
  - Monitoring Infinite Volumes
  - Configuring Infinite Volume and storage class thresholds
  - Configuring rules and data policy
- Support for monitoring node root volumes
- Enhancement in Snapshot copies monitoring, including computing reclaimable space and deleting Snapshot copies
- Annotations for storage objects
- Integrates with OnCommand Workflow Automation to execute workflows for storage classes and monitor SVMs with Infinite Volume that do not have storage classes

Related concepts

What you can do with Unified Manager on page 14

What you can do with Unified Manager

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

The key capabilities of Unified Manager include monitoring, alerting, availability and capacity management of clusters, protection capabilities, performance monitoring, configuration and management of Infinite Volumes, 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 also plan the storage requirements of your cluster objects using the information provided in the capacity and health charts for the respective cluster object.

Related concepts

What the Unified Manager server does on page 20

Understanding the workspace

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

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

Typical window layouts

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

Not every window contains every element in the following diagrams.

Object list window layout

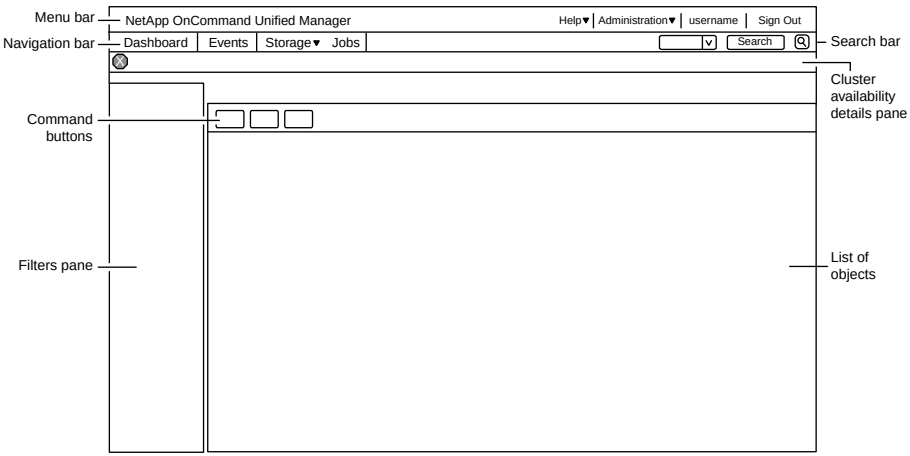

Details window layout

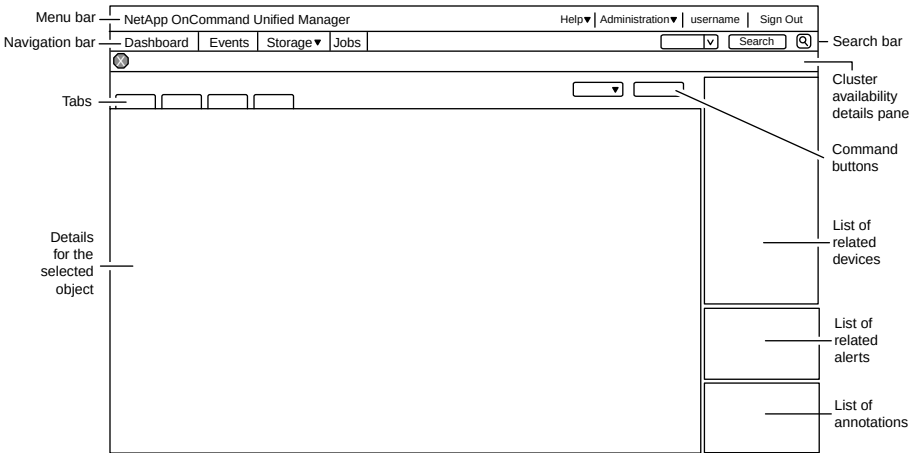

Related references

Window layout customization on page 16

Window layout customization

OnCommand Unified Manager enables you to customize the layout of information on the storage object pages. By customizing the windows, you can control which data is viewable or how it is displayed.

Sorting

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

When you click the column header, the sort arrow (and) appears for that column.

Filtering

You can apply filters to customize the display of information on the storage object pages so that only those entries that match the conditions provided are displayed. You can apply filters either from the Filters pane or on the columns.

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

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

type `vol?` to view the list of all volumes with the name containing "vol" followed by one more character—for example, vol1 or vol2. You can type `vol` to view the list of all volumes with the name starting with “vol”.

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

If you use the filter in the Available Data % column in the Aggregates page to filter data greater than 20.45%, actual capacity of 20.454% is displayed as 20.45%. Similarly, if you filter data less than 20.45%, actual capacity of 20.449% is displayed as 20.45%. For columns that display capacity data in percentage, you can view values up to four decimal points by moving your mouse pointer over the value displayed in the column.

Hiding or redisplaying the columns

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

Exporting data

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

Using the Unified Manager Help

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

About this task

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

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

Choices

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

  You can find more detailed information by expanding any portion of the Table of Contents in the navigation pane.
- To search the Help contents, click the **Search** tab in the navigation pane, type the word or series of words you want to find, and click **Go!**

- To print Help topics, click the printer icon.

Related tasks

Bookmarking your favorite topics on page 18

Bookmarking your favorite topics

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

Steps

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

Exporting data to CSV files

You can export data to a comma-separated values (`.csv`) file and use the exported data to build reports. For example, if there are 10 critical events that have not been resolved, you can export the data from the Events page to create a report and then take appropriate action.

Before you begin

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

About this task

You can export data from the following:

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

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

Steps

1. Perform one of the following actions:

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

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

   It might take a while for the data to be exported.
5. Click **OK**.

Related references

Volumes page on page 184
Aggregates page on page 175
Storage Virtual Machines page on page 158
Clusters page on page 144

What the Unified Manager server does

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

How the discovery process works

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

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

The following image illustrates the discovery process in OnCommand Unified Manager:

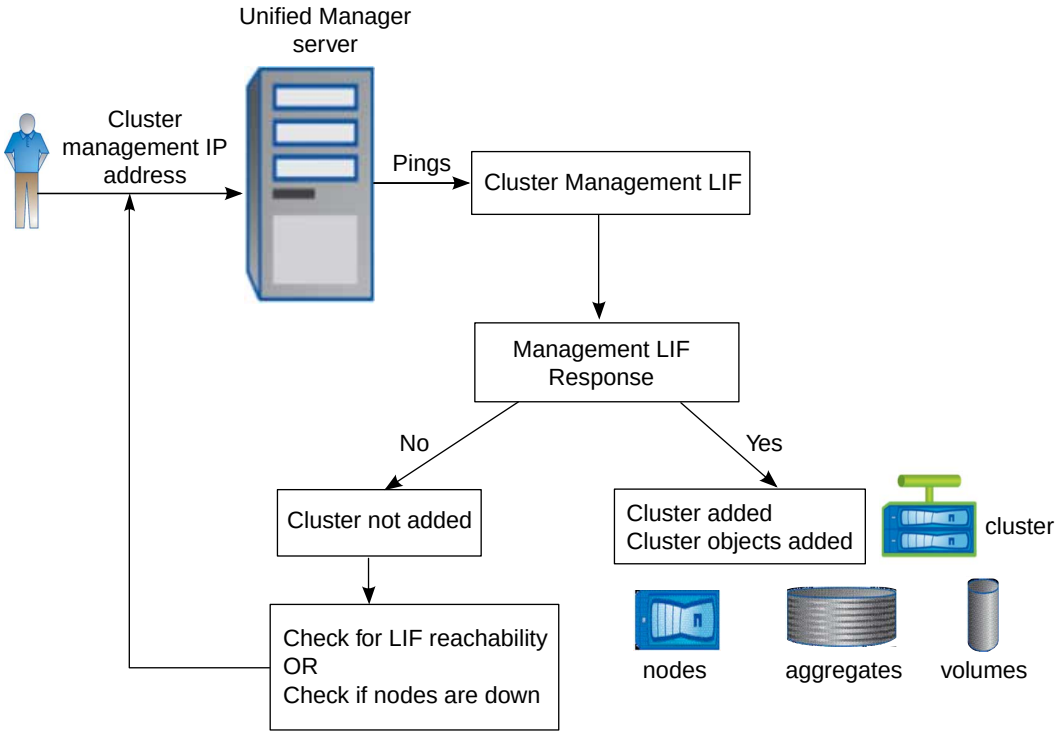

Searching for storage objects

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

Before you begin

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

Step

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

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

   #### Example

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

How timestamps work in Unified Manager

The timestamp displayed in Unified Manager is based on the time zone of the system on which you launch Unified Manager and not on the time zone that is configured for the Unified Manager server.

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

Getting started

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

Configuring your environment after deployment

After you deploy the Unified Manager virtual appliance, 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 deployed the virtual appliance and completed the Unified Manager initial setup.
- You must be logged in as the OnCommand Administrator role to perform this task.

About this task

After you complete the Unified Manager initial setup, you can add clusters. If you did not add clusters at that time, you must add them 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 prior to, or after, adding clusters.

Choices

- *Changing the Unified Manager host name* on page 23

  When you deployed Unified Manager, an SSL certificate was generated for HTTPS access. A host name was associated with the certificate, allowing you to use the host name to access the Unified Manager web UI. You might want to change this host name after deployment.
- *Configuring Unified Manager to send alert notifications* on page 24

  After the clusters have been added to Unified Manager, you can monitor them, but you cannot receive notifications about events in your cluster environment until you configure several options, such as the email address from which notifications are sent, the users to receive the alerts, and so forth. You might also want to modify the default threshold settings at which events are generated.
- *Adding clusters and viewing the discovery status* on page 120

  You must manually add clusters to Unified Manager before you can monitor them.

Changing the Unified Manager host name

When the virtual appliance is first deployed, the network host is assigned a name. You can change the host name after deployment. If you change the host name, you should also regenerate the HTTPS certificate.

Before you begin

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

About this task

You can use the host name (or the host IP address) to access the product 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 Workflow Automation. The host name is not updated automatically.

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

Steps

1. *Edit the network settings* on page 25

   You can change the host name from the Configure Network Settings dialog box, accessed from the Administration menu.

2. *Generate an HTTPS security certificate* on page 338

   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.

3. *View the HTTPS security certificate* on page 337

   You should verify that the correct information is displayed after generating a new security certificate, then restart Unified Manager.

4. *Restart the Unified Manager virtual machine* on page 338

   If you regenerate the HTTPS certificate, then you must restart the virtual machine.

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 be logged in as the OnCommand Administrator role to perform this task.

About this task

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

You can complete the following tasks to properly configure your environment and to add alerts.

Steps

1. *Configure notification settings* on page 49

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

2. *Enable remote authentication* on page 327

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

3. *Add authentication servers* on page 329

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

4. *Edit global threshold settings* on page 100

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

5. *Add users* on page 316

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

6. *Add alerts* on page 77

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

Customizing your environment

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

Enabling periodic AutoSupport

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

Before you begin

You must be logged in as the Maintenance User to perform this task.

Steps

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

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

Related tasks

Adding a user on page 316
Sending an on-demand AutoSupport message on page 344

Editing the network settings

You might want to edit network settings if an IP address changes due to the migration of a virtual machine (VM) to a different ESX server in a different domain, when maintenance is performed on your network equipment, if you switch from a DHCP to a static network configuration, or if you switch from a static network to a DHCP configuration.

Before you begin

- You might need one or more of the following: host name or FQDN, IP address, DHCP, network mask, gateway, primary and secondary DNS addresses, and search domains.

- If you are changing your network settings from DHCP-enabled to static network configuration, you should have done the following:
  - Ensured that the IP address and gateway are reachable
  - Ensured that the IP address does not contain a duplicate address
  - Verified that the primary and secondary DNS addresses are ready and available to send and receive network traffic
- You must be logged in as the OnCommand Administrator role to perform this task.

About this task

When you switch to a DHCP configuration, the previous host name is replaced by the name specified by your DHCP server.

The self-signed SSL certificate generated during deployment is associated with the host name (or FQDN) and the IP address. If you change either of these values and want to use that new host name or IP address to connect to Unified Manager, then you must generate a new certificate. The new certificate does not take effect until the Unified Manager virtual machine is restarted.

Steps

1. Click **Administration > Configure Network Settings**.
2. In the **Configure Network Settings** dialog box, modify the host and network settings, as required.

   Tip: You can enter multiple comma-separated values in the Secondary DNS Address and Search Domains fields.
3. Click **Update**.

After you finish

After you have modified the settings of your network configuration, you can use the updated configuration to access Unified Manager.

Related tasks

Adding a user on page 316

Editing the global setup options

You can use the Setup Options dialog box to configure Unified Manager global options. 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 options are not modified.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

- You can modify the following options:
  - Thresholds
    You can configure the global threshold values for aggregates, volumes, and relationships.
  - Management Server
    You can configure settings related to AutoSupport, HTTPS, and authentication.
  - General Settings
    You can configure settings related to events and notifications.
  - Quota Settings
    You can configure settings related to user and user group quota email notifications.
  - Add-ons
    You can configure settings for OnCommand Workflow Automation.
- You can also modify threshold settings for each object from the details page for that object.
- Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

Steps

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

Related tasks

Adding a user on page 316

Working with HTTPS security certificates

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

Before you begin

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

About this task

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

Choices

- *Generating an HTTPS security certificate* on page 338

  You might generate a new HTTPS security certificate for multiple reasons, including if you want to sign with a different Certificate Authority or if the current security certificate has expired. The new certificate replaces the existing certificate.
- *Downloading an HTTPS certificate signing request* on page 340

  You can download a certification request for the current HTTPS security certificate so that you can provide the file to a Certificate Authority to sign. A CA-signed certificate helps prevent man-in-the-middle attacks and provides better security protection than a self-signed certificate.
- *Installing an HTTPS security certificate* on page 341

  You can upload and install a security certificate after a Certificate Authority has signed and returned it. The file that you upload and install must be a signed version of the existing self-signed certificate. A CA-signed certificate helps prevent man-in-the middle attacks and provides better security protection than a self-signed certificate.
- *Viewing the HTTPS security certificate* on page 337

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

Related references

HTTPS Setup Options dialog box on page 342

Description of Setup options dialog boxes

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

Setup Options dialog box

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

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

- *Thresholds* on page 29
- *Management Server* on page 29
- *General Settings* on page 30
- *Quota Settings* on page 30
- *Add-ons* on page 30

- *Command buttons* on page 30

Thresholds

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

- Aggregates
  Specifies the global threshold values for monitored aggregates.
  For more detailed information, see the following topic:
  Aggregate Thresholds Setup Options dialog box on page 102
- Volumes
  Specifies the global threshold values for monitored volumes.
  For more detailed information, see the following topic:
  Volume Thresholds Setup Options dialog box on page 105
- Relationships
  Specifies the global threshold values for protection relationships.
  For more detailed information, see the following topic:
  Relationships Thresholds Setup Options dialog box on page 267

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

Management Server

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

- AutoSupport
  Enables you to view the AutoSupport description, enable periodic AutoSupport, or send an on-demand AutoSupport message.
  For more detailed information, see the following topic:
  AutoSupport Setup Options dialog box on page 32
- HTTPS
  Enables you to view the current security certificate, download a certificate signing request, to generate a new HTTPS certificate, or install a new HTTPS certificate.
  For more detailed information, see the following topic:
  HTTPS Setup Options dialog box on page 342
- Authentication
  Allows you to enable authentication, configure the settings to retrieve data from authentication servers, add or delete authentication servers, and test authentication.
  For more detailed information, see the following topic:
  Authentication Setup Options dialog box on page 332

General Settings

The General Settings area enables you to configure notification settings:

- Notification
  Enables you to specify an email address, SMTP server details, and SNMP trap details in order to receive notifications. SNMPv1 and SNMPv3 are supported.
  For more detailed information, see the following topic:
  Notification Setup Options dialog box on page 83

Quota Settings

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

- Email Address Rules
  Enables you to create rules to specify the email address of a user or user group to which emails are sent when there is a breach in the quotas.
  For more detailed information, see the following topic:
  Rules to Generate User and Group Quota Email Addresses dialog box on page 115
- Email Notification Format
  Enables you to create a notification format for the emails that are sent to a user or a user group when there is a breach in the quotas.
  For more detailed information, see the following topic:
  Email Notification Format dialog box on page 114

Add-ons

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

- Set Up OnCommand Workflow Automation
  Enables you to add, modify, or delete the OnCommand Workflow Automation settings.
  For more detailed information, see the following topic:
  Set Up OnCommand Workflow Automation on page 292

Command buttons

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

Restore to Factory Defaults	Enables you to restore the configuration settings to the factory default values.
Save	Saves the configuration settings for the selected option.
Save and Close	Saves the configuration settings for the selected option and closes the Setup Options dialog box.
Cancel	Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Editing the global setup options on page 26

Related references

Rules to Generate User and Group Quota Email Address dialog box on page 115
Email Notification Format dialog box on page 114

Configure Network Settings dialog box

You must configure the required network settings to connect to the management server. You can use the Configure Network Settings dialog box to modify the settings of your network configuration.

Host

The Host area provides the host name:

Host Name	Displays the host name of the system on which the management server is installed.

Network

The Network area provides information about the network, such as the IP address, network mask, and DNS information:

DHCP Enabled	Specifies if DHCP is enabled. **Note:** If DHCP is enabled, the system populates the IP address, network mask, and gateway fields with values from the network, and these fields appear dimmed.
IP Address	Specifies the IP address of the server.
Network Mask	Specifies the network mask.
Gateway	Specifies the IP address of the gateway.
Primary DNS Address	Specifies the IP address of the primary DNS server.
Secondary DNS Address	Specifies the IP address of the secondary DNS server.
Search Domains	Specifies the domain names (as comma-separated values) that are used by the DNS server to search for the host name.

Command buttons

The command buttons enable you to update or cancel the network configuration settings:

Update	Updates the network configuration settings.
Cancel	Discards the changes and closes the Configure Network Settings dialog box.

Related tasks

Editing the network settings on page 25

AutoSupport Setup Options dialog box

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

Information

You can perform the following operations:

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

Actions

You can perform the following operations:

Enable Periodic AutoSupport	Enables you to have specific, predefined messages to technical support for issue diagnosis and resolution periodically generated.
Generate and Send AutoSupport	Enables you to generate an on-demand message to send to technical support for any issues that have recently occurred.

Command buttons

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

Restore to Factory Defaults	Enables you to restore the configuration settings to the factory default values.
Save	Saves the configuration settings for the selected option.
Save and Close	Saves the configuration settings for the selected option and closes the Setup Options dialog box.
Cancel	Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Enabling periodic AutoSupport on page 25
Sending an on-demand AutoSupport message on page 344

Monitoring your system from the dashboard

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

Understanding the dashboard

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

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

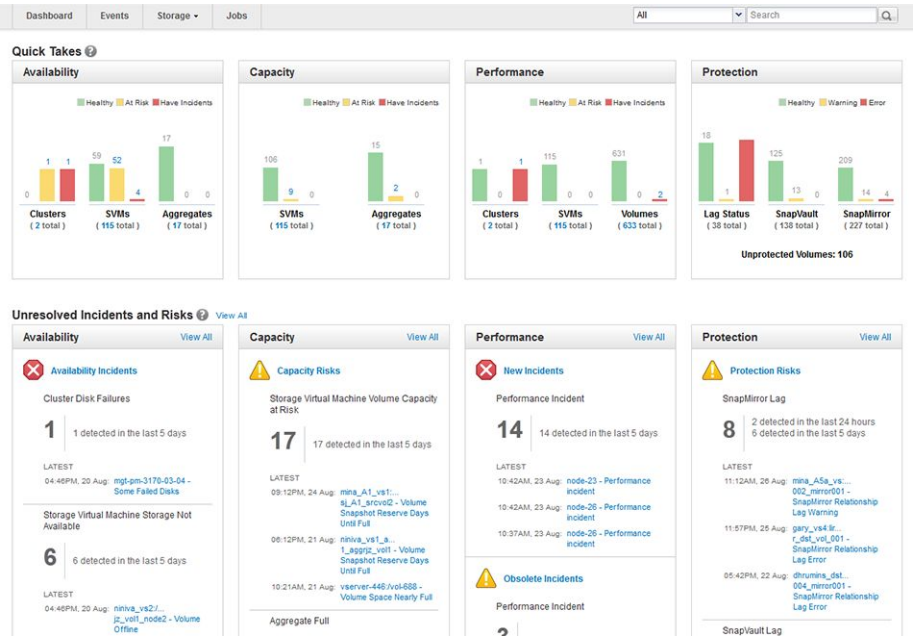

Quick Takes area

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

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

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

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

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

Unresolved Incidents and Risks area

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

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

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

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

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

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

Related references

Description of the Dashboard window on page 35

Description of the Dashboard window

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

Cluster availability details pane

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

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

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

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

Quick Takes area

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

Availability pane

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

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

Capacity pane

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

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

Performance pane

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

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

Protection pane

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

This pane also displays the total number of storage objects in each of the protection categories. Clicking the link for either the lag status, SnapVault, or SnapMirror category takes you to a filtered list of those object types in the Protection view of the Volumes page.

Unresolved Incidents and Risks area

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

Availability pane

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

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

Capacity pane

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

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

Performance pane

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

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

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

Protection pane

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

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

Monitoring performance

The performance monitoring support in the Unified Manager web UI enables you to track the I/O performance of your monitored storage and connect to related applications to diagnose storage performance issues.

Enabling performance monitoring

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

Creating a user with Event Publisher role privileges

To support a connection between a Performance Manager server and the Unified Manager server and the display of Performance Manager performance information in the Unified Manager web UI, you must first create a local user and assign to it the Event Publisher role.

Before you begin

You must be logged in to Unified Manager as the OnCommand Administrator to perform this task.

About this task

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

Steps

1. Click **Administration** > **Manage Users**.
2. In the **Manage Users** page, click **Add**.
3. In the **Add User** dialog box, select **Local User** for `type` and **Event Publisher** for `role` and enter the other required information.
4. Click **Add**.

Result

After you create a local user with the Event Publisher role, you can configure a connection between one or more Performance Manager servers and the Unified Manager server.

Related concepts

Purpose of a connection between Performance Manager and Unified Manager on page 44

Related tasks

Configuring a connection between a Performance Manager server and the Unified Manager server on page 39

Configuring a connection between a Performance Manager server and the Unified Manager server

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

Before you begin

- You must have created a local user with Event Publisher role privileges for the Unified Manager server in the connection you want to create.
- You must have a user ID authorized to log in to the maintenance console of the Performance Manager server for which you want to display performance data in the Unified Manager web UI.
- You must be prepared to specify the following information about the Unified Manager server:
  - Unified Manager server name or IP address
  - Unified Manager server default port 443
  - Event Publisher user name (the name of the local Unified Manager server user assigned Event Publisher role privileges)
  - Event Publisher password (the password of the local Unified Manager server user assigned Event Publisher role privileges)

About this task

You can configure connections between one Unified Manager server and multiple Performance Manager servers.

For each connection, complete the following actions:

Steps

1. Log in as the maintenance user to the maintenance console of the Performance Manager server for which you want to create the Unified Manager connection.
2. In the maintenance console, type the number of the menu option labeled "Unified Manager Server Connection" and then type the number of the menu option labeled "Add/Modify Unified Manager Server Connection."
3. When prompted, supply the requested Unified Manager server name or IP address and Unified Manager server port information.

The maintenance console checks the validity of the specified Unified Manager server name or IP address and Unified Manager server port, and, if necessary, prompts you to accept the Unified Manager server trust certificate to support the connection. The default Unified Manager server port 443 must be used.

4. When prompted, supply the requested Event Publisher user name and Event Publisher password and then confirm that the settings are correct.

Result

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

Related concepts

Purpose of a connection between Performance Manager and Unified Manager on page 44
Display of performance incidents in Unified Manager on page 44

Deleting a connection between a Performance Manager server and the Unified Manager server

To stop display in the Unified Manager web UI of performance issues discovered by a specific Performance Manager server, you must delete the connection between that server and the Unified Manager server in the Performance Manager maintenance console.

Before you begin

You must have a user ID authorized to log in to the maintenance console of the Performance Manager server whose performance data you want to stop displaying in the Unified Manager web UI.

Steps

1. Log in as the maintenance user to the maintenance console of the Performance Manager server whose Unified Manager connection you want to delete.
2. In the maintenance console, type the number of the menu option labeled “Unified Manager Server Connection” to display the **Unified Manager Connection Menu**, and then type the number of the menu option labeled “Delete Unified Manager Server Connection.”
3. When prompted to confirm that you want to continue the deletion, type the following:

   y

4. Press any key to continue.
5. To exit the maintenance console, press any key to continue, then type the following:

   x

Result

After the Unified Manager connection is deleted, notices of the performance incidents discovered by the disconnected Performance Manager server are no longer displayed in the Unified Manager web UI.

Viewing details of performance incidents listed in the Dashboard

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

Before you begin

- Unified Manager must be connected with one or more Performance Manager servers.
- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.
- You must be able to log in to the Performance Manager server associated with the performance incident that you want to diagnose.

Steps

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

   This action invokes Performance Manager to display in a new tab on your browser.

2. If prompted, log in to Performance Manager.

   Performance Manager displays the details page for the performance incident that you want to examine.

3. Follow the procedures described in the Performance Manager help to further determine the cause of the performance incident.

4. When you are finished, click the OnCommand Unified Manager browser tab to return to the Unified Manager web UI.

Related concepts

Understanding the dashboard on page 33
What performance incidents are on page 43
Display of performance incidents in Unified Manager on page 44

Viewing all performance incidents

To assess the overall state of your storage performance health, you can list all performance incidents on the Unified Manager Events page.

Before you begin

- Unified Manager must be connected with one or more Performance Manager servers.
- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.

Step

1. In the **Performance** pane of the Unified Manager **Dashboard** page, click **View All**.

   This action opens the Events page and displays all the New, Assigned, Resolved, and Obsolete performance incidents.

Related concepts

What performance incidents are on page 43
Display of performance incidents in Unified Manager on page 44

Related references

Events page on page 85

Viewing details of performance incidents listed on the Events page

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

Before you begin

- Unified Manager must be connected with one or more Performance Manager servers.
- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.
- You must be able to log in to the Performance Manager server associated with the performance incident that you want to diagnose.

Steps

1. Click **Events**.

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

   All the unresolved performance incidents are listed together for you to survey.

3. In the Name column, locate and click the hypertext link for the performance incident in question.

   This action displays a new tab on your browser for the display of Performance Manager.

4. If prompted, log in to Performance Manager.

   Performance Manager displays the details page for the performance incident that you want to examine.

5. Follow the procedures described in the Performance Manager help to further determine the cause of the performance incident.

6. When you have finished click the OnCommand Unified Manager browser tab to return to the Unified Manager web UI.

Related concepts

What performance incidents are on page 43
Display of performance incidents in Unified Manager on page 44

Related references

Events page on page 85

Understanding more about performance monitoring

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

What performance incidents are

Performance incidents are events generated by the performance monitoring application, OnCommand Performance Manager. These events indicate resource contention and I/O performance issues on aggregates and cluster nodes that require your attention.

Performance incidents occur when workload activities cause unacceptable increases in the time it takes for a volume on a cluster to respond to read and write requests from applications.

Purpose of a connection between Performance Manager and Unified Manager

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

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

Connections between multiple Performance Manager servers and Unified Manager

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

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

Display of performance incidents in Unified Manager

If the Unified Manager server is connected to a Performance Manager server, then performance incidents generated by Performance Manager are displayed in the Unified Manager web UI. This enables you to monitor the performance health of your managed storage using the same application and from the same dashboard and event windows that you also use to monitor the availability, capacity, and protection health of your managed storage.

Where performance incidents are displayed

Within the Unified Manager web UI, the performance incidents are graphed in the Quick Takes area of the Dashboard page, listed in the Unresolved Incidents and Risks area of the Dashboard page, and listed on the Events page.

What information the performance graph displays

The performance graph in the Quick Takes area of the Dashboard page graphs the number of clusters, Storage Virtual Machines (SVMs), and volumes monitored by the Performance Manager server and the Unified Manager server that are "Healthy" or "Have Incidents." On the performance graph, the “Have Incidents” bars show the number of clusters, SVMs, and volumes whose I/O times are adversely affected by contention issues on the cluster, cluster node, or aggregate resources.

What information the performance incident listings display

All performance events displayed in the Unified Manager web UI are uniformly worded "Performance incident" and are listed with the name of the cluster, cluster node, or aggregate object whose resource contention issues are causing the performance incident. Their display in the Unified Manager web UI merely alerts you that performance issues for the resource in question have occurred.

To view the details of a particular performance incident that is listed in the Unified Manager web UI and diagnose its cause, you must click its hypertext links, which open the appropriate pages in Performance Manager.

What states of performance incidents are displayed

The states of performance incidents displayed in the Unified Manager web UI are periodically refreshed at intervals between 5-15 minutes. New performance incidents or performance incidents that flag conditions that still persist are displayed in the “New” state. Performance incidents that flag conditions that no longer exist are displayed in the “Obsolete” state. Performance incidents that have been acknowledged or assigned are displayed in the “Acknowledged” or “Assigned” state. Performance incidents that have been resolved are counted as Obsolete on the Dashboard page and displayed as Resolved on the Events page.

Workloads, cluster components, and performance

The Performance Manager application, whether functioning as a standalone application or as an application connected to Unified Manager, helps you diagnose and address performance issues by identifying and analyzing instances of workloads exceeding what a cluster component can supply, accommodate, or allow.

What a workload is

A workload is the storage activity (the processing bandwidth or throughput bandwidth) that is required by common storage functions such as deduplication, RAID reconstruction, Snapshot copy backup, WAFL management, volume moves, or flexible volume maintenance and read/write activity.

Workloads that are monitored by Performance Manager include the following:

FlexVol volume workload	The storage activity associated with maintaining a volume and its read/write activity. Multiple FlexVol volumes constitute multiple workloads that must be supported by the shared cluster components on which the volumes reside or through which volume data is throughput. Because FlexVol volumes are created and configured by storage administrators, volume workloads are also referred to as *user-defined workloads*.

Deduplication workload	The storage activity associated with data deduplication activity on a storage cluster. Because deduplication is a system-defined feature, deduplication workloads are referred to as *system workloads*.
RAID reconstruction workload	The storage activity associated with RAID reconstruction of stored data that might be temporarily lost due to a single or double disk failure in a storage array. RAID reconstruction workloads are system workloads.
Snapshot copy workload	The storage activity associated with the automatic backup imaging of data in Snapshot copies. Snapshot copy workloads are system workloads.
WAFL consistency point (CP) workload	The storage activity associated with maintaining the WAFL consistency point feature, which enables fast restart and recovery of storage cluster nodes in the event of improper shutdown. WAFL consistency point workloads are system workloads.
WAFL scan	The storage activity associated with execution of the WAFL scan operation. WAFL scan operations are system workloads.
Volume moves workload	The storage activity associated with flexible volume migration. Volume move workloads are system workloads.
Other internal activity	The storage activity associated with other system-related operations.

What a cluster component is

A cluster component is a physical or logical component of the cluster that supplies, accommodates, or constrains the storage activity that workloads use to execute their functions.

Performance Manager monitors certain categories of cluster components and generates performance incidents when the workload storage activity demands exceed what those components can supply, accommodate, or allow, thus causing workload I/O response times to lengthen beyond an acceptable threshold.

Performance Manager monitors and, if necessary, generates performance incidents for the following categories of cluster components:

Network	The I/O processing from clients connected to the cluster. If network issues occur outside of the cluster (for example, high network traffic or over-utilized clients), thus causing slow workload I/O response times, Performance Manager generates a performance incident related to this.
Policy Group limit	The Storage Quality of Service (QoS) policy group of which the workload is a member. If the processing demand of the workloads in the policy group exceed the policy limit for I/O operations, the workloads are throttled, and Performance Manager generates a performance incident.
Network Processing	The software component in the cluster involved with I/O processing between the network protocols and the cluster. If high response time is detected during

	network processing, Performance Manager generates a performance incident with the name of the cluster node that is handling that network processing.
Cluster Interconnect	The cables and adapters with which clustered nodes are physically connected. Performance Manager generates a performance incident if it detects contention in the interconnect.
Data Processing	The software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains a volume workload. If Performance Manager detects a high response time between the cluster and the storage aggregate, it generates a performance incident with the name of the cluster node that is handling the data processing. The I/O requests might or might not be from other nodes in the cluster.
Aggregate	The storage aggregate that contains the volume workload. If the processing demands of one workload increase the response time of other workloads on the aggregate, Performance Manager generates a performance incident with the name of the aggregate in contention.

How performance incidents are discovered and analyzed

The performance incidents that you see displayed in the Unified Manager web UI are actually discovered and analyzed by the Performance Manager server or servers to which it is connected.

Performance incident discovery

Performance Manager monitors the I/O response time and operations for volumes on a cluster. When other workloads overuse a cluster component, for example, the component is in contention and is unable to perform at an optimal level to meet the combined workload demands.

Incident analysis

When the response time for a volume workload crosses the performance threshold, Performance Manager analyzes the event by performing the following actions:

- Identifies the “victim” workloads (those whose performance has decreased)
- Identifies the cluster component in contention
- Identifies the “bully” workloads (those that are overusing the cluster component and causing it to be in contention)
- Ranks the workloads involved to determine which bullies have the highest change in usage of the component and which victims are the most impacted.

An incident might occur for only a brief moment and then correct itself after the component it is using is no longer in contention. An incident might also recur. When an incident is resolved, it remains listed in Performance Manager as part of the record of past performance issues for a volume. Each incident has a unique ID that identifies the incident type and the volume, cluster, and cluster components involved.

Managing events and alerts

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

What events are

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

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

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

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

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

Related concepts

What alerts are on page 77

Related references

Description of event severity types on page 55
Description of event impact levels on page 56
Description of event impact areas on page 56

Configuring event settings

You can specify the number of days an event is retained in the Unified Manager server before it is automatically deleted. Only events that are resolved, obsolete, or of type Information are deleted.

You can also specify the frequency with which these events are deleted or you can also manually delete the events.

Before you begin

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

About this task

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

Steps

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

Related tasks

Adding a user on page 316

Configuring notification settings

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

Before you begin

The following information must be available:

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

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

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

3. In the **Notification Setup Options** dialog box, configure the appropriate settings.

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

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

Related tasks

Adding a user on page 316
Adding an alert on page 77

Viewing event details

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

Before you begin

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

About this task

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

Steps

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

   The event details are displayed in the Event details page.

Related tasks

Adding a user on page 316
Assigning events on page 52
Adding an alert on page 77
Acknowledging and resolving events on page 54
Adding and reviewing notes about an event on page 53

Disabling or enabling events

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

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

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

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

Steps

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

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

Assigning events

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

Before you begin

- The user's name and email ID must be configured correctly.
- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.

Steps

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

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

Related tasks

Adding a user on page 316

Viewing unassigned events

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

Before you begin

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

Steps

1. Click **Events**.

   By default, New and Acknowledged events are displayed on the Events page.

2. From the **Filters** pane, select the **Unassigned** filter option in the **Assigned To** area.

Related tasks

Adding a user on page 316
Assigning events on page 52

Adding and reviewing notes about an event

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

Before you begin

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

Steps

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

Related tasks

Adding a user on page 316

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 be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.

About this task

You can acknowledge and resolve multiple events simultaneously.

Steps

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

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

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

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

Related tasks

Adding a user on page 316
Adding and reviewing notes about an event on page 53

Understanding more about events

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

Event state definitions

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

The different event states are as follows:

New	The state of a new event.
Acknowledged	The state of an event when you have acknowledged it.
Resolved	The state of an event when it is marked as resolved.
Obsolete	The state of an event when it is automatically corrected or when the cause of the event is no longer valid. **Note:** You cannot acknowledge or resolve an obsolete event.

Example of different states of an event

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

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

Description of event severity types

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

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

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

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

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

Related references

List of events and severity types on page 57

Description of event impact levels

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

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

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

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

Description of event impact areas

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

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

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

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

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

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

How object status is computed

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

Related concepts

Event state definitions on page 55

Related tasks

Acknowledging and resolving events on page 54

List of events and severity types

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

Aggregate events

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

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Aggregate Offline (ocumEvtAggregateStateOffline)	Incident	Aggregate	Critical

Event name (Trap name)	Impact level	Source type	Severity
Aggregate Failed (ocumEvtAggregateStateFailed)	Incident	Aggregate	Critical
Aggregate Restricted (ocumEvtAggregateStateRestricted)	Risk	Aggregate	Warning
Aggregate Reconstructing (ocumEvtAggregateRaidStateReconstructing)	Risk	Aggregate	Warning
Aggregate Degraded (ocumEvtAggregateRaidStateDegraded)	Risk	Aggregate	Warning

Impact area: capacity

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

Impact area: configuration

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

Cluster events

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

Impact area: availability

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

Impact area: configuration

Event name **(Trap name)**	**Impact level**	**Source type**	**Severity**
Node Added (Not applicable)	Event	Cluster	Information
Node Removed (Not applicable)	Event	Cluster	Information

Event name (Trap name)	Impact level	Source type	Severity
Cluster Removed (Not applicable)	Event	Cluster	Information
Cluster Add Failed (Not applicable)	Event	Cluster	Error
Cluster Name Changed (Not applicable)	Event	Cluster	Information

Port events

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

Impact area: availability

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

Disks events

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

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Some Unassigned Disks (ocumEvtClusterUnassignedDisksSome)	Risk	Cluster	Warning
Some Failed Disks (ocumEvtDisksSomeFailed)	Incident	Cluster	Critical

Enclosures events

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

Impact area: availability

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

Impact area: configuration

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

Fans events

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

Impact area: availability

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

Flash card events

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

Impact area: availability

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

Inodes events

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

Impact area: capacity

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

Logical interface events

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

Impact area: availability

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

Impact area: configuration

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

LUN events

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

Impact area: availability

Event name (Trap name)	**Impact level**	**Source type**	**Severity**
LUN Offline (ocumEvtLunOffline)	Incident	LUN	Critical
Single Active Path To Access LUN (ocumEvtLunSingleActivePath)	Risk	LUN	Warning

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

Impact area: capacity

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

Node events

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

Impact area: availability

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

Impact area: configuration

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

NVRAM battery events

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

Impact area: availability

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

Power supplies events

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

Impact area: availability

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

Protection events

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

Impact area: protection

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

Qtree events

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

Impact area: capacity

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

Service processor events

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

Impact area: availability

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

SnapMirror relationship events

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

Impact area: protection

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

Event name (Trap name)	Impact level	Source type	Severity
SnapMirror Relationship Resync Failed (ocumEvtSnapmirrorRelationshipResyncFailed)	Risk	SnapMirror relationship	Error
SnapMirror Relationship Deleted ocumEvtSnapmirrorRelationshipDeleted	Risk	SnapMirror relationship	Warning

SnapVault relationship events

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

Impact area: protection

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

Storage failover settings events

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

Impact area: availability

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

Storage services events

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

Impact area: configuration

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

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

Impact area: protection

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

Storage shelf events

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

Impact area: availability

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

User and group quota events

User and group quota events provide you with information about the capacity of the user and user group quota as well as the file and disk limits so that you can monitor for potential problems. Events

are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

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

Volume events

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

Impact area: availability

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

Event name (Trap name)	Impact level	Source type	Severity
Volume Remounted (Not applicable)	Event	Volume	Information
Volume Junction Path Inactive (ocumEvtVolumeJunctionPathInactive)	Risk	Volume	Warning

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Thin-Provisioned Volume Space At Risk (ocumThinProvisionVolumeSpaceAtRisk)	Risk	Volume	Warning
Volume Space Full (ocumEvtVolumeFull)	Risk	Volume	Error
Volume Space Nearly Full (ocumEvtVolumeNearlyFull)	Risk	Volume	Warning
Volume Snapshot Reserve Space Full (ocumEvtSnapshotFull)	Risk	Volume	Warning
Too Many Snapshot Copies (ocumEvtSnapshotTooMany) This event is applicable only for volumes in a cluster running Data ONTAP 8.2 or later.	Risk	Volume	Error
Volume Qtree Quota Overcommitted (ocumEvtVolumeQtreeQuotaOvercommitted)	Risk	Volume	Error
Volume Qtree Quota Nearly Overcommitted (ocumEvtVolumeQtreeQuotaAlmostOvercommitted)	Risk	Volume	Warning
Volume Growth Rate Abnormal (ocumEvtVolumeGrowthRateAbnormal)	Risk	Volume	Warning
Volume Days Until Full (ocumEvtVolumeDaysUntilFullSoon)	Risk	Volume	Error
Volume Space Guarantee Disabled (Not applicable)	Event	Volume	Information

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

Impact area: configuration

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

Volume move status events

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

Impact area: capacity

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

Event name (Trap name)	Impact level	Source type	Severity
Volume Move - Cutover Deferred (ocumEvtVolumeMoveCutoverDeferred)	Risk	Volume	Warning

SVM events

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

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
SVM CIFS Service Down (ocumEvtVserverCifsServiceStatusDown)	Incident	SVM	Critical
SVM FC/FCoE Service Down (ocumEvtVserverFcServiceStatusDown)	Incident	SVM	Critical
SVM iSCSI Service Down (ocumEvtVserverIscsiServiceStatusDown)	Incident	SVM	Critical
SVM NFS Service Down (ocumEvtVserverNfsServiceStatusDown)	Incident	SVM	Critical
SVM CIFS Service Not Configured (Not applicable)	Event	SVM	Information
SVM FC/FCoE Service Not Configured (Not applicable)	Event	SVM	Information
SVM iSCSI Service Not Configured (Not applicable)	Event	SVM	Information
SVM NFS Service Not Configured (Not applicable)	Event	SVM	Information
SVM Stopped (ocumEvtVserverDown)	Risk	SVM	Warning

Event name (Trap name)	Impact level	Source type	Severity
SVM with Infinite Volume Storage Not Available (ocumEvtVserverStorageNotAvailable)	Incident	SVM SVMs with Infinite Volume	Critical
SVM with Infinite Volume Storage Partially Available (ocumEvtVserverStoragePartiallyAvailable)	Risk	SVM SVMs with Infinite Volume	Error
SVM with Infinite Volume Namespace Mirror Constituents Having Availability Issues (ocumEvtVserverNsMirrorAvailabilityHavingIssues)	Risk	SVM SVMs with Infinite Volume	Warning

Impact area: capacity

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

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

Impact area: configuration

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

SVM storage class events

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

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

Impact area: availability

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

Impact area: capacity

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

Event name (Trap name)	Impact level	Source type	Severity
SVM Storage Class Snapshot Usage Limit Exceeded (ocumEvtVserverStorageClassSnapshotUsageExceeded)	Risk	Storage class	Warning

What alerts are

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

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

Related concepts

What events are on page 48

Adding an alert

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

Before you begin

- You must have configured notification settings such as the email address, SMTP server, and SNMP trap host so that the Unified Manager server can use these settings to send notifications to users when an event is generated.
- The following information must be available: resources and events for which you want to trigger the alert, and user names or email addresses of users you want to notify.
- You must be logged in as the OnCommand Administrator role to perform this task..

About this task

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

Steps

1. Click **Administration > Manage Alerts**.
2. In the **Alerts** page, click **Add**.
3. In the **Add Alert** dialog box, perform the following steps:
   a) Click **Name** and enter a name and description for the alert.

b) Click **Resources** and select the resources to be included 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 you specify, the list of available resources displays only those resources that match the filter rule.

Note: 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.

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

c) 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.

d) Click **Recipients** and select the users that you want to notify when the alert is generated and the notification frequency.

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 have modified the email address of the selected user from the Manage Users page, the modified email address is not updated for the selected user.

You can also choose to notify users through SNMP traps.

4. Click **Save**.

Example for adding an alert

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

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

Perform the following steps in the Add Alert dialog box:

1. Click **Name** and enter `Test` in the **Alert Name** field.

2. Click **Resources** and in the Include tab, select Volumes from the drop-down list.

a. Enter `abc` in the **Name contains** field to display the volumes whose name contains abc.

b. Select **<<All Volumes whose name contains 'abc'>>** from the Available Resources area and move it to the Selected Resources area.

c. Click **Exclude** and enter `xyz` in the **Name contains** field and then click **Add**.

3. Click **Events** and select Critical from the **Event Severity** field.
4. Select **All Critical Events** from the Matching Events area and move it to the Selected Events area.
5. Click **Recipients** and enter `sample@domain.com` in the **Alert these users** field.
6. Select **Remind every 15 minutes** to set the frequency to notify the user every 15 minutes.
   You can configure an alert to repeatedly send notification to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.
7. Click **Save**.

Related concepts

What alerts are on page 77
Guidelines for adding alerts on page 79

Related tasks

Adding a user on page 316
Configuring notification settings on page 49

Guidelines for adding alerts

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

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

- Alert description
  You should provide a description for the alert so that it helps you track your alerts effectively.
- Resources
  You should decide which physical or logical resource requires an alert. You can include and exclude resources, as required. For example, if you want to closely monitor your aggregates by configuring an alert, you must select the required aggregates from the list of resources.
- Event severity
  You should decide if an event of a specified severity type (Critical, Error, Warning) should trigger the alert and, if so, which severity type.
- Event name
  If you add an alert based on the type of event generated, you should decide which events require an alert
- Recipients

You must provide the user names and email addresses of the users who receive the notification. You can also specify an SNMP trap as a mode of notification.

- Notification frequency

  You can configure an alert to repeatedly send notification to the recipients for a specified time. You should determine the time from which the event notification is active for the alert. If you want the event notification to be repeated until the event is acknowledged, you should determine how often you want the notification to be repeated.

Testing an alert

You can test an alert to verify that you have configured it correctly. A test alert is sent and you can verify if you have received the notification.

Before you begin

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

Steps

1. Click **Administration** > **Alerts**.
2. In the **Alerts** page, select the alert that you want to test and click **Test**.

   A test alert is sent to the email addresses that are specified during alert creation.

Related tasks

Adding a user on page 316
Configuring notification settings on page 49

Viewing alerts

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

Before you begin

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

Step

1. Click **Administration** > **Manage Alerts**.

   The list of alerts is displayed in the Alerts page.

Related tasks

Adding a user on page 316

Editing an alert

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

Before you begin

You must have the role of OnCommand Administrator to edit an alert.

Steps

1. Click **Administration** > **Manage Alerts**.
2. In the **Alerts** page, select the alert you want to edit and click **Edit**.
3. In the **Edit Alert** dialog box, edit the required information.
4. Click **Save**.

Related concepts

What alerts are on page 77

Related tasks

Adding a user on page 316

Deleting alerts

You can delete one or more alerts when they are no longer needed. For example, you might want to delete an alert that was created for a particular resource because the resource might no longer be monitored by the Unified Manager server.

Before you begin

You must have the role of OnCommand Administrator to delete alerts.

Steps

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

Related tasks

Adding a user on page 316

Description of Event windows and dialog boxes

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

Event Retention Settings

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

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Event Settings

You can configure the following options:

Delete Information, Resolved, and Obsolete Events Older Than	Enables you to specify the retention period after which information, resolved, and obsolete events are removed from the management server. The default value is 180 days. Retaining the events for more than 180 days affects the performance and is not recommended. The lower limit for the event retention period is 7 days, although there is no upper limit.
Delete Schedule	Enables you to specify the frequency at which all the information, resolved, and obsolete events that have exceeded their age limit are automatically deleted from the management server. The possible values are Daily, Weekly, or Monthly. The default value is Daily.
Delete Now	Enables you to manually delete all the information, resolved, and obsolete events that have exceeded their specified retention period.

Command buttons

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

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

Related tasks

Configuring event settings on page 48

Notification Setup Options dialog box

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

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Email

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

From Address	Specifies the email address from which the alert notification is sent.

SMTP Server

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

Host Name	Specifies the host name of your SMTP host server, which is used to send the alert notification to the specified recipients.
User Name	Specifies the SMTP user name. SMTP user name is required only when the SMTPAUTH is enabled in the SMTP server.
Password	Specifies the SMTP password. SMTP user name is required only when the SMTPAUTH is enabled in the SMTP server.
Default Port	Specifies the port that is used by the SMTP host server to send alert notification. The default value is 25.
Use SSL	Selecting this box provides secure communication between the SMTP server and the management server.

SNMP

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

Version	Specifies the SNMP version you want to use depending on the type of security you require. Options include Version 1, Version 3, Version 3 with Authentication, and Version 3 with Authentication and Encryption. The default value is Version 1.
Trap Destination Host	Specifies the host that receives the SNMP traps that are sent by the management server.
Outbound Trap Port	Specifies the port through which the SNMP server receives the traps that are sent by the management server. The default value is 162.
Community	Specifies the community name that is used by the SNMP server to authenticate the traps that are sent from the management server. The community name is only available with SNMP Version 1.
Engine ID	Specifies the unique identifier of the SNMP agent and is automatically generated by the management server. Engine ID is available with SNMP Version 3, SNMP Version 3 with Authentication, and SNMP Version 3 with Authentication and Encryption.
User Name	Specifies the SNMP user name. User name is available with SNMP Version 3, SNMP Version 3 with Authentication, and SNMP Version 3 with Authentication and Encryption.
Authentication Protocol	Specifies the protocol used to authenticate a user. Protocol options include MD5 and SHA. MD5 is the default value. Authentication protocol is available with SNMP Version 3 with Authentication and SNMP Version 3 with Authentication and Encryption.
Authentication Password	Specifies the password used when authenticating a user. Authentication password is available with SNMP Version 3 with Authentication and SNMP Version 3 with Authentication and Encryption.
Privacy Protocol	Specifies the privacy protocol used to encrypt SNMP messages. Protocol options include AES 128 and DES. The default value is AES 128. Privacy protocol is available with SNMP Version 3 with Authentication and Encryption.
Privacy Password	Specifies the password when using privacy protocol. Privacy password is available with SNMP Version 3 with Authentication and Encryption.

Command buttons

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

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

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

Related tasks

Configuring notification settings on page 49

Events page

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

- *Command buttons* on page 85
- *Events list* on page 86
- *Filters pane* on page 87

Command buttons

The command buttons enable you to perform the following tasks:

Assign To	**Me**	Assigns the event to you. When you assign an event to yourself, your user name and the time when you assigned the event is added in the events list for the selected events.
	Another user	Displays the Assign Owner dialog box, which enables you to assign or reassign the event to other users. When you assign an event to a user, the user's name and the time when the event was assigned are added in the events list for the selected events. **Note:** You can also unassign events by leaving the ownership field blank.
Acknowledge	Acknowledges the selected events. When you acknowledge an event, your user name and the time when you acknowledged the event are added in the events list for the selected events. When you acknowledge an event, you are responsible for managing that event.	
Mark As Resolved	Enables you to change the event state to resolved.	

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

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

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

Events list

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

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

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

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

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

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

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

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

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

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

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

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

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

Acknowledged Time	Displays the time that has elapsed since the event was acknowledged. If the time elapsed exceeds a week, the timestamp when the event was acknowledged is displayed. By default, this column is hidden.
Resolved By	Displays the name of the user who resolved the event. The field is blank if the event is not resolved. By default, this column is hidden.
Resolved Time	Displays the time that has elapsed since the event was resolved. If the time elapsed exceeds a week, the timestamp when the event was resolved is displayed. By default, this column is hidden.
Obsoleted Time	Displays the time when the state of the event became Obsolete. By default, this column is hidden.

Filters pane

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

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

Related tasks

Viewing event details on page 50
Assigning events on page 52
Viewing unassigned events on page 53
Acknowledging and resolving events on page 54
Adding an alert on page 77
Exporting data to CSV files on page 18

Event details page

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

- *Command buttons* on page 88
- *Summary area* on page 88
- *Notes and Updates area* on page 89
- *Possible Causes* on page 89
- *Resources that Might be Impacted* on page 90
- *Suggested Corrective Actions area* on page 90

Command buttons

The command buttons enable you to perform the following tasks:

Assign To	**Me**	Assigns the event to you.
	Another user	Opens the Assign Owner dialog box, which enables you to assign or reassign the event to other users. When you assign an event to a user, the user's name and the time when the event was assigned are added in the events list for the selected events. **Note:** You can also unassign events by leaving the ownership field blank.
Acknowledge	Acknowledges the selected events so that you do not continue to receive repeat alert notifications.	
Mark As Resolved	Enables you to change the event state to resolved.	
Add Alert	Displays the Add Alert dialog box, which enables you to add an alert for the selected event.	
View Events	Navigates to the Events page.	

Summary area

You can view the following event details:

Severity	Displays the severity of the event. The event severity types are Critical ( ), Error ( ), Warning ( ), and Information ( ).
State	Displays the event state: New, Acknowledged, Resolved, or Obsolete.
Impact Level	Displays whether the event is categorized as an incident, risk, or an informational event.
Impact Area	Displays whether the event is a capacity, availability, protection, performance, or configuration related event.
Obsoleted Cause	Displays the reason the event is now obsolete.
Source	Displays the full name of the object along with the type of object with which the event is associated. The value is displayed as “Unknown” when Data ONTAP does not provide a valid user name because of SecD errors.
Source Type	

	Displays the object type (for example, Storage Virtual Machine (SVM), volume, or qtree) with which the event is associated.
Acknowledged By	Displays the name of the person who acknowledged the event and the time when 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 when 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 selected event, based on the recent timestamp. You can also view the time when the information was added.

Post	Enables you to display the information that you added.

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 the resource.

The area is displayed only for the Volume Offline event, the Volume Restricted event, and the Thin-Provisioned Volume Space At Risk event.

Resources that Might be Impacted area

Displays the resources that might be impacted because of the availability issues in your volume if the Volume Offline event or the Volume Restricted event is generated and the capacity issues in the aggregate if the Thin-Provisioned Volume Space At Risk event is generated. You can click the name of the resource to view more details about the impacted resource.

The area is displayed only for the Volume Offline event, the Volume Restricted event, and the Thin-Provisioned Volume Space At Risk event.

Suggested Corrective Action area

Displays the actions that you can perform to address the capacity issues of your volume.

The area is displayed only for the Volume Space Nearly Full event and the Volume Space Full event.

Related tasks

Viewing event details on page 50
Assigning events on page 52
Acknowledging and resolving events on page 54
Adding an alert on page 77
Adding and reviewing notes about an event on page 53

Manage Events page

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

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

- *Command buttons* on page 90
- *List view* on page 91

Command buttons

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

Disable	Displays the Disable Events dialog box, which you can use to disable events.
Enable	Enables events.
Event Retention Settings	Displays the Event Retention Settings dialog box, which enables you to specify the retention period after which the information, resolved, and obsolete events are removed from the management server. The default retention value is 180 days.

List view

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

Event	Displays the name of the event that is disabled.
Severity	Displays the severity of the event. The severity can be Critical, Error, Warning, or Information.
Source Type	Displays the source type for which the event is generated.

Disable Events dialog box

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

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

- *Event properties area* on page 91
- *Command buttons* on page 91

Event Properties area

The Event Properties area area specifies the following event properties:

Event Severity	Enables you to select events based on the severity type, which can be Critical, Error, Warning, or Information.
Event Name Contains	Enables you to filter events whose name contains the specified characters.
Matching events	Displays the list of events matching the event severity type and the text string you specify.
Disable events	Displays the list of events that you have selected for disabling. The severity of the event is also displayed along with the event name.

Command buttons

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

Save and close	Disables the event type and closes the Disable Events dialog box.
Cancel	Discards the changes and closes the Disable Events dialog box.

Description of Alert windows and dialog boxes

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

Alerts page

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

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

- *Command buttons* on page 92
- *List view* on page 92
- *Details area* on page 93

Command buttons

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

Add	Displays the Add Alert dialog box, which enables you to add new alerts.
Edit	Displays the Edit Alert dialog box, which enables you to edit the selected alert.
Delete	Deletes the selected alerts.
Enable	Enables the selected alerts to send notifications.
Disable	Disables the selected alerts when you want to temporarily stop sending notifications.
Test	Tests the selected alert to verify its configuration after adding or editing an alert.

List view

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

Status	Specifies whether an alert is enabled ( ) or disabled ( ).
Alert	Displays the name of the alert.
Description	Displays a description for the alert.
Notification Method	Displays the notification method that is selected for the alert. You can notify users through email or SNMP traps.

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

Details area

The details area provides more information about the selected alert.

Alert Name	Displays the name of the alert.
Alert Description	Displays a description for the alert.
Events	Displays the events for which you want to trigger the alert.
Resources	Displays the resources for which you want to trigger the alert.
Includes	Displays the group of resources for which you want to trigger the alert.
Excludes	Displays the group of resources for which you do not want to trigger alert.
Notification Method	Displays the notification method for the alert.
Notification Frequency	Displays the frequency with which the management server continues to send notifications until the event is acknowledged, resolved, or moved to the obsolete state.
Email Recipients	Displays the email addresses of users who receive the notification.

Related tasks

Adding an alert on page 77
Editing an alert on page 81
Testing an alert on page 80
Deleting alerts on page 81
Viewing alerts on page 80

Add Alert dialog box

You can create alerts to notify you when a particular event is generated. You can create alerts for a single resource, set of resources, and events of a particular severity type. You can also specify the notification method and frequency of the alerts.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Name

This area enables you to specify a name and description for the alert:

Alert Name	Enables you to specify an alert name.

Alert Description	Enables you to specify a description for the alert.

Resources

This area enables you to select an individual resource or group the resources based on a dynamic rule for which you want to trigger the alert. A dynamic rule is the set of resources filtered based on the text string you specify. You can search for resources by selecting a resource type from the drop-down list or you can specify the exact resource name to display a specific resource.

Note: If you are creating an alert from any of the storage object details pages, the storage object is automatically included in the alert.

Include Enables you to include the resources for which you want to trigger alerts. You can specify a text string to group resources that match the string and select this group to be included in the alert. For example, you can group all volumes whose name contains the “abc” string.

Exclude Enables you to exclude resources for which you do not want to trigger alerts. For example, you can exclude all volumes whose name contains the “xyz” string.

Note: The Exclude tab is displayed only when you select all resources of a particular resource type—for example, <<All Volumes>> or <<All Volumes whose name contains '*xyz*'>>.

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

Events

This area enables you to select the events for which you want to create the alerts. You can create alerts for events based on a particular severity or for a set of events.

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

Event Severity	Enables you to select events based on the severity type, which can be Critical, Error, or Warning.
Event Name Contains	Enables you to select events whose name contains the specified characters.

Recipients

This area enables you to specify the users that you want to notify when an alert is triggered. You can also specify the notification method and the frequency of notification.

Alert these users	Enables you to specify the email address or user name of the user to receive notifications.

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 have modified the email address of the selected user from the Manage Users page, the modified email address is not updated for the selected user.

Notification Frequency	Enables you to specify the frequency with which the management server sends notifications until the event is acknowledged, resolved, or moved to the obsolete state. You can choose the following notification methods: • Notify only once • Notify at a specified frequency • Notify at a specified frequency within the specified time range
Issue SNMP trap	Checking this box enables you to specify whether SNMP traps should be sent to the globally configured SNMP host.

Command buttons

The command buttons enable you to perform the following tasks:

Save	Creates an alert and closes the Add Alert dialog box.
Cancel	Discards the changes and closes the Add Alert dialog box.

Related tasks

Adding an alert on page 77

Edit Alert dialog box

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

Name

This area enables you to edit the name and description for the alert.

Alert Name	Enables you to edit the alert name.
Alert Description	Enables you to specify a description for the alert.
Alert State	Enables you to enable or disable the alert.

Resources

This area enables you to select an individual resource or group the resources based on a dynamic rule for which you want to trigger the alert. You can search for resources by selecting a resource type from the drop-down list or you can specify the exact resource name to display a specific resource.

Include Enables you to include the resources for which you want to trigger alerts. You can specify a text string to group resources that match the string and select this group to be included in the alert. For example, you can group all volumes whose name contains the “vol0” string.

Exclude Enables you to exclude resources for which you do not want to trigger alerts. For example, you can exclude all volumes whose name contains the “xyz” string.

> **Note:** The Exclude tab is displayed only when you select all resources of a particular resource type—for example, <<All Volumes>> or <<All Volumes whose name contains '*xyz*'>>.

Events

This area enables you to select the events for which you want to trigger the alerts. You can trigger an alert for events based on a particular severity or for a set of events.

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

Event Name Contains Enables you to select events whose name contains the specified characters.

Recipients

This area enables you to specify the notification method and the frequency of notification.

Alert these users Enables you to edit the email address or user name, or specify a new email address or user name to receive notifications.

Notification Frequency Enables you to edit the frequency with which the management server sends notifications until the event is acknowledged, resolved, or moved to the obsolete state.

You can choose the following notification methods:

- Notify only once
- Notify at a specified frequency
- Notify at a specified frequency within the specified time range

Issue SNMP trap Enables you to specify whether SNMP traps should be sent to the globally configured SNMP host.

Command buttons

The command buttons enable you to perform the following tasks:

Save	Saves the changes and closes the Edit Alert dialog box.
Cancel	Discards the changes and closes the Edit Alert dialog box.

Related tasks

Editing an alert on page 81

Managing thresholds

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

What storage capacity thresholds are

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

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

Configuring global aggregate threshold values

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

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

- Most of the threshold options have default values for better monitoring. However, you can change the values to suit the requirements of your environment.
- The threshold values are not applicable to the root aggregate of the node.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Thresholds > Aggregates**.
3. Configure the appropriate threshold values for capacity, growth, and Snapshot copies.
4. Click **Save and Close**.

Related concepts

Understanding capacity events and thresholds for node root aggregates on page 119

Related tasks

Adding a user on page 316

Configuring global volume threshold values

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

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

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

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Thresholds > Volumes**.
3. Configure the appropriate threshold values for capacity, Snapshot copies, quotas, volume growth, and inodes.
4. Click **Save and Close**.

Related tasks

Adding a user on page 316

Editing unmanaged relationship lag threshold settings

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

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Thresholds > Relationships**.
3. In the **Lag** area of the **Relationships Thresholds Setup Options** dialog box, increase or decrease the warning or error lag time percentage as needed.
4. Click **Save and Close**.

Related tasks

Adding a user on page 316

Editing global threshold settings

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

About this task

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

Global threshold settings are accessible from the Setup Options dialog box. You can also modify threshold settings for individual objects, from the inventory page or the details page for that object.

Choices

- *Configuring global aggregate threshold values* on page 98

  You can edit the threshold settings for capacity, growth, and Snapshot copies for all aggregates to track any threshold breach.
- *Configuring global volume threshold values* on page 99

  You can edit the threshold settings for capacity, Snapshot copies, quotas, volume growth, overwrite reserve space, and inodes for all volumes to track any threshold breach.
- *Editing unmanaged relationship lag thresholds* on page 99

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

Editing aggregate threshold settings

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

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

- Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.
- The threshold values are not applicable to the root aggregate of the node.

Steps

1. Click **Storage** > **Aggregates**.

   You can also edit the thresholds from the Administration menu, Setup Options dialog box.
2. In the **Aggregates** page, select one or more aggregates and click **Edit Thresholds**.
3. In the **Edit Aggregate Thresholds** dialog box, perform the appropriate action:

If you are...	Then...
Editing the threshold settings of one aggregate	Edit the appropriate settings.
Editing the threshold settings of multiple aggregates	Select the appropriate check box and edit the settings.

4. Click **Save and Close**.

Related concepts

Understanding capacity events and thresholds for node root aggregates on page 119

Related tasks

Adding a user on page 316

Editing volume threshold settings

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

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

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

Steps

1. Click **Storage** > **Volumes**.

   You can also edit the thresholds from the Administration menu, Setup Options dialog box.
2. On the **Volumes** page, select one or more volumes and click **Edit Thresholds**.
3. In the **Edit Volume Thresholds** dialog box, perform the appropriate action:

If you are...	Then...
Editing the threshold settings of one volume	Edit the appropriate settings.
Editing the threshold settings of multiple volumes	Select the appropriate check box and edit the settings.

4. Click **Save and Close**.

Related tasks

Adding a user on page 316

Description of thresholds dialog boxes

You can use the appropriate Setup Options dialog box to configure global threshold values for aggregates and volumes.

Aggregate Thresholds Setup Options dialog box

The Aggregate Thresholds Setup Options dialog box enables you to configure global threshold values for monitored aggregates. You can set thresholds for individual aggregates or for all the

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

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Events are generated when a threshold is breached. You can take corrective actions for such events.

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

You can set thresholds for the following: capacity, aggregate growth, and aggregate Snapshot copies.

- *Capacity* on page 103
- *Growth* on page 104
- *Snapshot Copies* on page 104
- *Command buttons* on page 105

Capacity area

The Capacity area enables you to set the following aggregate capacity threshold conditions:

Space Nearly Full	Specifies the percentage at which an aggregate is considered to be nearly full: • Default value: 80 percent The value for this threshold must be lower than the value for the Aggregate Full threshold for the management server to generate an event. • Event generated: Aggregate Nearly Full • Event severity: Warning
Space Full	Specifies the percentage at which an aggregate is considered full: • Default value: 90 percent • Event generated: Aggregate Full • Event severity: Error
Nearly Overcommitted	Specifies the percentage at which an aggregate is considered to be nearly overcommitted: • Default value: 95 percent The value for this threshold must be lower than the value for the Aggregate Overcommitted Full threshold for the management server to generate an event. • Event generated: Aggregate Nearly Overcommitted • Event severity: Warning
Overcommitted	Specifies the percentage at which an aggregate is considered to be overcommitted:

- Default value: 100 percent
- Event generated: Aggregate Overcommitted
- Event severity: Error

Days Until Full

Specifies the number of days remaining before the aggregate reaches full capacity:

- Default value: 7
- Event generated: Aggregate Days Until Full
- Event severity: Error

Growth area

The Growth area enables you to set the following threshold conditions for aggregate growth:

Growth Rate

Specifies the percentage at which an aggregate's growth rate is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event:

- Default value: 1 percent
- Event generated: Aggregate Growth Rate Abnormal
- Event severity: Warning

Growth Rate Sensitivity

Specifies the factor that is applied to the standard deviation of an aggregate's growth rate. If the growth rate exceeds the factored standard deviation, an Aggregate Growth Rate Abnormal event is generated.

A lower value for growth rate sensitivity indicates that the aggregate is highly sensitive to changes in the growth rate. The range for the growth rate sensitivity is 1 through 5.

- Default value: 2

Note: If you modify the growth rate sensitivity for aggregates at the global threshold level, the change is also applied to the growth rate sensitivity for volumes at the global threshold level.

Snapshot copies area

The Snapshot copies area enables you to set the following Snapshot reserve threshold conditions:

Snapshot Reserve Full

Specifies the percentage at which an aggregate has consumed all the space reserved for Snapshot copies:

- Default value: 90 percent
- Event generated: Aggregate Snapshot Reserve Full
- Event severity: Warning

Command buttons

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

Restore to Factory Defaults	Enables you to restore the configuration settings to the factory default values.
Save	Saves the configuration settings for the selected option.
Save and Close	Saves the configuration settings for the selected option and closes the Setup Options dialog box.
Cancel	Cancels the recent changes and closes the Setup Options dialog box.

Related concepts

Understanding capacity events and thresholds for node root aggregates on page 119

Related tasks

Configuring global aggregate threshold values on page 98

Volume Thresholds Setup Options dialog box

The Volume Thresholds Setup Options dialog box enables you to configure global threshold values for monitored volumes. You can set thresholds for individual volumes or for all the volumes globally. When you configure the options globally, the default values of the objects are modified. However, if the default values have been changed at the object level, the global values are not modified.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Events are generated when a threshold is breached. You can take corrective actions for such events.

You can set thresholds for the following: capacity, volume Snapshot copies, quotas, volume growth, and inodes.

- *Capacity* on page 105
- *Snapshot Copies* on page 106
- *Qtree Quota* on page 107
- *Growth* on page 107
- *Inodes* on page 107
- *Command buttons* on page 108

Capacity area

The Capacity area enables you to set the following volume capacity threshold conditions:

Space Nearly Full

Specifies the percentage at which a volume is considered to be nearly full:

- Default value: 80 percent

  The value for this threshold must be lower than the value for the Volume Full threshold in order for the management server to generate an event.
- Event generated: Volume Nearly Full
- Event severity: Warning

Space Full

Specifies the percentage at which a volume is considered full:

- Default value: 90 percent
- Event generated: Volume Full
- Event Severity: Error

Days Until Full

Specifies the number of days remaining before the volume reaches full capacity:

- Default value: 7
- Event generated: Volume Days Until Full
- Event severity: Error

Snapshot copies area

The Snapshot copies area enables you to set the following threshold conditions for the Snapshot copies in the volume:

Snapshot Reserve Full

Specifies the percentage at which the space reserved for Snapshot copies is considered full:

- Default value: 90 percent
- Event generated: Volume Snapshot Reserve Full
- Event severity: Error

Days Until Full

Specifies the number of days remaining before the space reserved for Snapshot copies reaches full capacity:

- Default value: 7
- Event generated: Volume Snapshot Reserve Days Until Full
- Event severity: Error

Count

Specifies the number of Snapshot copies that can be created on a volume before the system generates the Too Many Snapshot Copies event:

- Default value: 250
- Event generated: Too Many Snapshot Copies
- Event severity: Error

Note: This field is applicable only for volumes in a cluster running Data ONTAP 8.2 or later.

Qtree Quota area

The Qtree Quota area enables you to set the following volume quota threshold conditions:

Nearly Overcommitted

Specifies the percentage at which a volume is considered to be nearly overcommitted by qtree quotas:

- Default value: 95 percent
- Event generated: Volume Qtree Quota Nearly Overcommitted
- Event severity: Warning

Overcommitted

Specifies the percentage at which a volume is considered to be overcommitted by qtree quotas:

- Default value: 100 percent
- Event generated: Volume Qtree Quota Overcommitted
- Event severity: Error

Growth

The Growth area enables you to set the following threshold conditions for volume growth:

Growth Rate

Specifies the percentage at which a volume's growth rate is considered to be normal before the system generates a Volume Growth Rate Abnormal event:

- Default value: 1 percent
- Event generated: Volume Growth Rate Abnormal
- Event severity: Warning

Growth Rate Sensitivity

Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated.

A lower value for growth rate sensitivity indicates that the volume is highly sensitive to changes in the growth rate. The range for the growth rate sensitivity is 1 through 5.

- Default value: 2

Note: If you modify the growth rate sensitivity for volumes at the global threshold level, the change is also applied to the growth rate sensitivity for aggregates at the global threshold level.

Inodes

The Inodes area enables you to set the following threshold conditions for inodes:

Nearly Full

Specifies the percentage at which a volume is considered to have consumed most of its inodes:

- Default value: 80 percent
- Event generated: Inodes Nearly Full
- Event severity: Warning

Full Specifies the percentage at which a volume is considered to have consumed all of its inodes:

- Default value: 90 percent
- Event generated: Inodes Full
- Event severity: Error

Command buttons

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

Restore to Factory Defaults	Enables you to restore the configuration settings to the factory default values.
Save	Saves the configuration settings for the selected option.
Save and Close	Saves the configuration settings for the selected option and closes the Setup Options dialog box.
Cancel	Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Configuring global volume threshold values on page 99

Managing quotas

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

What quota limits are

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

By default, the Unified Manager server sends a notification email to users who have crossed the quota soft limit or have reached the quota hard limit, and user quota events are generated. The OnCommand Administrator can configure alerts that notify the specified recipients of the user or user group quota events.

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

Viewing user and user group quotas

The Storage Virtual Machine details page displays information about the user and user group quotas that are configured on the SVM. You can view the name of the user or user group, soft and hard limits set on the disks and files, used disk and file space, and disk threshold value.

Before you begin

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

Steps

1. Click **Storage > Storage Virtual Machines**.
2. Select the SVM for which you want to view the user and user group quota details.
3. Click **User and Group Quotas**.

Related concepts

Overview of the quota process on page 113

Related tasks

Adding a user on page 316

Creating rules to generate email addresses

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

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have reviewed the *guidelines for creating rules* on page 115.

About this task

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

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

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

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

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Quota Settings > Email Address Rules**.
3. In the **Rules to Generate User and Group Quota Email Address** dialog box, enter the rule based on your criteria.
4. Click **Validate** to validate the syntax of the rule.

   An error message is displayed if the syntax of the rule is incorrect. You must correct the syntax and click **Validate** again.
5. Click **Save and Close**.
6. Verify that the email address you created is displayed in the **User and Group Quotas** tab of the **Storage Virtual Machine details** page.

Related tasks

Editing user and group quota email addresses on page 111
Creating an email notification format for user and user group quotas on page 111

Creating an email notification format for user and user group quotas

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

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Quota Settings > Email Notification Format**.
3. Enter or modify the details in the **From**, **Subject**, and **Email Details** fields.
4. Click **Preview** to preview the email notification.

   The Notification Preview dialog box is displayed with the email notification.
5. Click **Close**.
6. Modify the content, if required.
7. Click **Save**.

Related tasks

Creating rules to generate email addresses on page 110

Editing user and group quota email addresses

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

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

- You must have reviewed the *guidelines for creating rules* on page 115.

About this task

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

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. Select the SVM for which you want to view the user and user group quota details.
3. Click **User and Group Quotas**, and then click **Edit Email Address**.
4. In the **Edit Email Address** dialog box, perform the appropriate action:

If...	Then...
You want notifications to be sent to the email address generated by the rules specified	**a.** Delete the email address in the **Email Address** field. **b.** Click **Save**. **c.** Refresh the browser by pressing F5 to reload the Storage Virtual Machine details page. The email address generated by the specified rule is displayed in the **Email Address** field.
You want notifications to be sent to a specified email address	**a.** Modify the email address in the **Email Address** field. **b.** Click **Save**. The rules to generate the user and group quota email addresses are no longer applicable to the quota.

Related tasks

Creating rules to generate email addresses on page 110

Understanding more about quotas

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

Overview of the quota process

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

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

- Whether the hard limit will be reached

  If yes, the write operation will fail when the hard limit is reached and the hard quota notification is sent.
- Whether the soft limit will be breached

  If yes, the write operation will succeed when the soft limit is breached and the soft quota notification is sent.
- Whether a write operation will not exceed the soft limit

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

Related concepts

About quotas on page 113

Why you use quotas on page 113

About quotas

Quotas provide a way to restrict or track the disk space and number of files used by a user, group, or qtree.

Why you use quotas

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

You specify a quota for the following reasons:

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

Description of quotas dialog boxes

You can use the appropriate Setup Options dialog box to configure the format of the email notification that is sent when a quota-related issue occurs and to configure rules to specify email addresses based on the user quota.

Email Notification Format dialog box

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

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

From
Displays the email address from which the email is sent, which can be modified. By default, this is the email address that is specified under Notification in the General Settings area of the Setup Options dialog box.

Subject
Displays the subject of the notification email.

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

Valid keywords are as follows:

- $EVENT_NAME
  Specifies the event name that caused the email notification.
- $QUOTA_TARGET
  Specifies the qtree or volume on which the quota is applicable.
- $QUOTA_USED_PERCENT
  Specifies the percentage of disk hard limit, disk soft limit, file hard limit, or file soft limit that is used by the user or user group.
- $QUOTA_LIMIT
  Specifies one of the following:
  - Disk soft limit or file soft limit that is breached by the user or user group and the User or Group Quota Disk Space Soft Limit Breached event or the User or Group Quota File Count Soft Limit Breached event is generated.
  - Disk hard limit or file hard limit that is reached by the user or user group and the User or Group Quota Disk Space Hard Limit Reached event or the User or Group Quota File Count Hard Limit Reached event is generated.
- $QUOTA_USED

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

- $QUOTA_USER

  Specifies the user or user group name.

Command buttons

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

Preview	Displays a preview of the notification email.
Restore to Factory Defaults	Enables you to restore the notification format to the factory default values.
Save	Saves the changes made to the notification format.
Save and Close	Saves the changes made to the notification format and closes the Email Notification Format dialog box.
Cancel	Cancels the recent changes and closes the Email Notification Format dialog box.

Related references

Email Notification Format dialog box on page 114
Setup Options dialog box on page 28
Rules to Generate User and Group Quota Email Address dialog box on page 115

Rules to Generate User and Group Quota Email Address dialog box

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

Rules area

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

How you define rules

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

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

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

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

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

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

How you add comments

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

Rules syntax

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

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

Command buttons

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

Validate	Validates the syntax of the created rule. If there are errors during validation, the rule that generates the error is displayed along with an error message.

Save	Validates the syntax of the rule and saves the rule if there are no errors. If there are errors during validation, the rule that generates the error is displayed along with an error message.
Save and Close	Validates the syntax of the rule, saves the rule if there are no errors, and closes the Rules to Generate User and Group Quota Email Address dialog box.
Cancel	Cancels the recent changes and closes the Rules to Generate User and Group Quota Email Address dialog box.

Related tasks

Creating rules to generate email addresses on page 110

Related references

Setup Options dialog box on page 28
Email Notification Format dialog box on page 114

Managing and monitoring clusters and cluster objects

Unified Manager uses periodic API queries and a data collection engine to collect data from the clusters. By adding clusters to the Unified Manager database, you can monitor and manage these clusters for any availability and capacity risks.

Understanding cluster monitoring

You can add clusters to the Unified Manager database to monitor clusters for availability, capacity, and other details, such as CPU usage, interface statistics, free disk space, qtree usage, and chassis environmental.

Events are generated if the status is abnormal or when a predefined threshold is breached. If configured to do so, Unified Manager sends a notification to a specified recipient when an event triggers an alert.

The following flowchart illustrates the Unified Manager monitoring process:

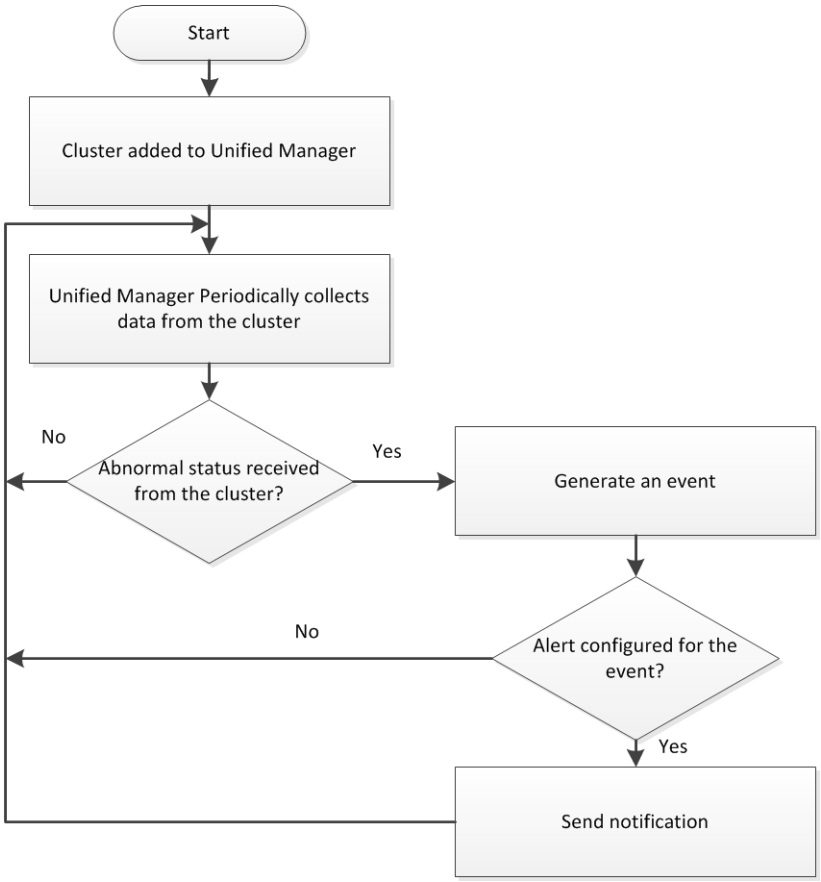

Understanding node root volumes

You can monitor the node root volume using Unified Manager. The best practice is that the node root volume should have sufficient capacity to prevent the node from going down.

When the used capacity of the node root volume exceeds 80 percent of the total node root volume capacity, the Node Root Volume Space Nearly Full event is generated. You can configure an alert for the event to get a notification. You can take appropriate actions to prevent the node from going down by using either OnCommand System Manager or the Data ONTAP CLI.

Understanding capacity events and thresholds for node root aggregates

You can monitor the node root aggregate by using Unified Manager. The best practice is to thickly provision the root volume in the root aggregate to prevent the node from halting.

By default, capacity events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to the node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are

modified by the technical support representative, the threshold values are applied to the node root aggregate.

You can take appropriate actions to prevent the node from halting by using either OnCommand System Manager or the Data ONTAP CLI.

Adding clusters

You can add an existing cluster to Unified Manager to monitor the cluster and obtain information such as its health, capacity, and configuration. You can also view the cluster discovery status from the Data Sources page.

Before you begin

- The following information must be available:
  - 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. This host name must resolve to the cluster-management IP address.

    The cluster-management IP address must be the cluster-management LIF of the admin Storage Virtual Machine (SVM). If you use a node-management LIF, the operation fails.
  - Data ONTAP administrator user name and password
  - Type of protocol (HTTP or HTTPS) that can be configured on the cluster and the port number of the cluster
- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Storage** > **Clusters**.
2. From the **Clusters** page, click **Add**.
3. In the **Add Cluster** dialog box, specify the values, such as the host name or IP address of the cluster, user name, password, protocol for communication, and port number.

   By default, the HTTPS protocol is selected.
4. Click **Add**.
5. If HTTPS is selected, perform the following steps:

   a) In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information of the cluster.

   b) Click **Yes**.

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

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

For more information, see *KB article 1014389 How to renew an SSL certificate in clustered Data ONTAP* (login required).

6. Optional: View the cluster discovery status by performing the following steps:
   a) Click the **Data Sources** link from the discovery status message displayed in the **Clusters** page.
   b) Review the cluster discovery status from the **Data Sources** page.

Result

The cluster is added to the Unified Manager database after the default monitoring interval of approximately 15 minutes.

Related tasks

Adding a user on page 316
Removing a cluster in case of discovery failure on page 121
Restarting cluster discovery on page 122

Removing a cluster in case of discovery failure

If the discovery of a cluster fails, you can remove that cluster by using the Data Sources page.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

The **Remove** button in the Data Sources page is enabled only when a selected cluster's discovery fails.

Steps

1. Click **Administration** > **Manage Data Sources**.
2. Select the cluster whose discovery has failed.

   You can use the filter on the State column to list the clusters whose discovery has failed.
3. Click **Remove**.

Related tasks

Adding clusters on page 120

Related references

Data Sources page on page 205

Restarting cluster discovery

You can restart the discovery of a cluster in case of any discovery failure by using the Data Sources page. For example, you might want to restart cluster discovery when there are network issues while adding the cluster to the Unified Manager database.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Administration** > **Manage Data Sources**.
2. Select the cluster that is not discovered.

   You can use the filter on the State column to list the clusters that are not discovered.
3. Click **Rediscover**.

Related references

Data Sources page on page 205

Editing clusters

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

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Storage** > **Clusters**.
2. On the **Clusters** page, select the cluster you want to edit and click **Edit**.
3. In the **Edit Cluster** dialog box, modify the values as required.
4. If you selected the HTTPS protocol, perform the following steps:

a) In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information of the cluster.
b) Click **Yes**.

5. Click **Save**.

Related tasks

Adding a user on page 316

Removing clusters

You can remove a cluster from Unified Manager by using the Clusters page. For example, you can remove a cluster when you want to decommission a storage system.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Storage** > **Clusters**.
2. On the **Clusters** page, select the cluster that you want to remove and click **Remove**.
3. Click **Yes** to confirm the remove request.

Result

The cluster, its storage objects along with the history, and all associated events are removed, and the cluster is no longer monitored by Unified Manager.

Related tasks

Adding a user on page 316

Rediscovering updated clusters manually

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

If Unified Manager is paired with OnCommand Workflow Automation, this operation also triggers the reacquisition of WFA cached data.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. From the **Clusters** page, click the cluster you want to rediscover.

   The Cluster details page is displayed.

2. Click **Actions** > **Rediscover**.

   The message “Cluster Rediscover Job Started” is displayed, along with a job link at the top of the Cluster details page.

3. Click **View details** in the job message to track the progress of the cluster rediscovery operation.

   The job task details for the cluster rediscovery job are displayed.

4. Click **Refresh** to update the job status.

   Cluster rediscovery is complete when all job tasks are finished.

Viewing the cluster list and details

You can use the Clusters page to view your inventory of clusters. You can also view details such as the cluster health, capacity, configuration, LIFs, and nodes in that cluster by using the Cluster details page.

Before you begin

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

About this task

The details in the Clusters page and the Cluster details page help you plan your storage. For example, before provisioning a new aggregate, you can select a specific cluster from the Clusters page and obtain capacity details to determine if the cluster has the required space.

Steps

1. Click **Storage** > **Clusters**.

2. View the complete details of the cluster by clicking the cluster name.

Related tasks

Adding a user on page 316
Exporting data to CSV files on page 18
Adding an alert on page 77
Editing clusters on page 122

Related references

Clusters page on page 144
Cluster details page on page 146

Viewing the node list and details

You can use the Nodes page to view the list of nodes in your clusters. You can use the Cluster details page to view detailed information about nodes that are part of the cluster that is monitored.

Before you begin

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

About this task

You can view details such as the node state, cluster that contains the node, aggregate capacity details (used and total), and raw capacity details (usable, spare, and total). You can also obtain information about HA pairs, disks shelves, and ports.

Steps

1. Click **Storage** > **Clusters**.
2. On the **Clusters** page, click the cluster name whose node details you want to view.
3. On the **Cluster details** page, click the **Nodes** tab.

   The left pane displays the list of HA pairs. By default, the HA Details tab is open, which displays HA state details and events related to the selected HA pair.
4. To view other details about the node, perform the appropriate action:

To view...	Click the...
Details about the disk shelves	**Disk Shelves** tab.
Port-related information	**Ports** tab.

Related tasks

Adding a user on page 316

Viewing the cluster list and details on page 124
Exporting data to CSV files on page 18
Adding an alert on page 77
Editing clusters on page 122

Related references

Nodes page on page 156

Viewing the SVM list and details

From the Storage Virtual Machines page, you can monitor your inventory of Storage Virtual Machines (SVMs). You can use the Storage Virtual Machine details page to view detailed information about SVMs that are monitored.

Before you begin

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

About this task

You can view SVM details, such as the capacity, efficiency, and configuration of an SVM. You can also view information about the related devices and related alerts for that SVM.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. Choose one of the following ways to view the SVM details:
   - To view minimal details, position the cursor over the SVM name.
   - To view the complete details, click the SVM name.

     You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. Optional: View the objects related to the SVM by clicking **View Related** in the minimal details dialog box.

Related tasks

Adding a user on page 316
Exporting data to CSV files on page 18
Adding an alert on page 77
Viewing the details of SVMs with Infinite Volume on page 294

Related references

Storage Virtual Machines page on page 158
Storage Virtual Machine details page on page 160

Viewing the aggregate list and details

From the Aggregates page, you can monitor your inventory of aggregates. You can use the Aggregate details page to view detailed information about aggregates that are monitored.

Before you begin

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

About this task

You can view details such as aggregate capacity and configuration, and disk information from the Aggregate details page. You can use these details before you modify the threshold settings if required.

Steps

1. Click **Storage** > **Aggregates**.
2. Choose one of the following ways to view the aggregate details:
   - To view minimal details, position the cursor over the aggregate name.
   - To view the complete details, click the aggregate name.
     You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. Optional: View the objects related to the aggregate by clicking **View Related** from the minimal details dialog box.

Related tasks

Adding a user on page 316
Exporting data to CSV files on page 18
Adding an alert on page 77
Editing aggregate threshold settings on page 101

Related references

Aggregates page on page 175
Aggregate details page on page 178

Viewing the volume list and details

From the Volumes page, you can monitor your inventory of volumes. You can use the Volume details page to view detailed information about volumes that are monitored, including the capacity, efficiency, and configuration of the volumes. You can also view information about the related devices and related alerts for a specific volume.

Before you begin

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

Steps

1. Click **Storage** > **Volumes**.
2. Choose one of the following ways to view the volume details:
   - To view minimal details, position the cursor over the volume name.
   - To view the complete details, click the volume name.
     You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. Optional: View the objects related to the volume by clicking **View Related** from the minimal details dialog box.

Related tasks

Adding a user on page 316
Exporting data to CSV files on page 18
Adding an alert on page 77
Editing volume threshold settings on page 102

Related references

Volumes page on page 184
Volume details page on page 190

Viewing the CIFS shares

You can use the Storage Virtual Machine details page to view detailed information about the CIFS share hosted by the Storage Virtual Machine (SVM). You can view details such as the share name, junction path, containing objects, security settings, and export policies defined for the share.

Before you begin

- CIFS license must be enabled on the cluster.
- LIFs serving the CIFS shares must be configured.
- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. Select the SVM for which you want to view the CIFS share details.
3. Click **CIFS Shares**.

Related concepts

About CIFS and SMB on page 143

Related tasks

Troubleshooting access to CIFS shares on page 347

Viewing the data sources

You can use the Data Sources page to view detailed information about the data sources that are added to the Unified Manager database. You can view details such as the discovery status, data source name and type, supported discovery operations, operation state, operation start and end time, and description of the operation.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Step

1. Click **Administration > Manage Data Sources**.

   Discovery details about the data sources are displayed in the Data Sources page.

Viewing the list of Snapshot copies

You can view the list of Snapshot copies for a selected volume. You can use the list of Snapshot copies to calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted, and you can delete the Snapshot copies if required.

Before you begin

- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.
- The volume containing the Snapshot copies must be online.

Steps

1. Click **Storage** > **Volumes**.
2. In the **Volumes** page, select the appropriate volume that contains the Snapshot copies you want to view.
3. In the **Capacity** tab of the **Volume details** page, click the link next to **Snapshot Copies**.

   For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. For volumes in a cluster running Data ONTAP 8.1.x, you should click the **View** link.

   The list of Snapshot copies is displayed.

Deleting Snapshot copies

You can delete a Snapshot copy to conserve space or to free disk space, or you can delete the Snapshot copy if it is no longer required.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

The volume must be online.

To delete a Snapshot copy that is busy or locked, you must have released the Snapshot copy from the application that was using it.

About this task

- You cannot delete the base Snapshot copy in a parent volume if a FlexClone volume is using that Snapshot copy.

The base Snapshot copy is the Snapshot copy that is used to create the FlexClone volume and displays the status `Busy` and Application Dependency as `Busy,Vclone` in the parent volume.

- You cannot delete a locked Snapshot copy that is used in a SnapMirror relationship.

  The Snapshot copy is locked and is required for the next update.

Steps

1. Click **Storage** > **Volumes**.
2. In the **Volumes** page, select the volume that contains the Snapshot copies you want to delete.
3. In the **Capacity** tab of the **Volume details** page, click the link next to **Snapshot Copies**.

   For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. For volumes in a cluster running Data ONTAP 8.1.x, you should click the **View** link.

   The list of Snapshot copies on the volume is displayed.
4. Select the Snapshot copies you want to delete.
5. Click **Delete Selected**.

Calculating reclaimable space for Snapshot copies

You can calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted.

Before you begin

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

The volume must be online.

Steps

1. Click **Storage** > **Volumes**.
2. In the **Volumes** page, select the appropriate volume for which you want calculate the reclaimable space.
3. In the **Capacity** tab of the **Volume details** page, click the link next to **Snapshot Copies**.

   For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. For volumes in a cluster running Data ONTAP 8.1.x, you should click the **View** link.

   The list of Snapshot copies on the volume is displayed.
4. Select the Snapshot copies for which you want to calculate the reclaimable space.

5. Click **Calculate**.

   The reclaimable space (in percentage, and KB, MB, GB, and so on) on the volume is displayed.

6. To recalculate the reclaimable space, select the required Snapshot copies and click **Recalculate**.

Understanding clusters and cluster objects

You should understand some of the basic concepts of clusters and cluster objects to monitor and manage your clusters.

What a cluster is

You can group pairs of nodes together to form a scalable cluster. Creating a cluster enables the nodes to pool their resources and distribute work across the cluster, while presenting administrators with a single entity to manage. Clustering also enables continuous service to end users if individual nodes go offline.

The maximum number of nodes within a cluster depends on the platform model and licensed protocols. For details about cluster size limits, see the *Hardware Universe* at *hwu.netapp.com*.

Each node in the cluster can view and manage the same volumes as any other node in the cluster. The total file-system namespace, which comprises all of the volumes and their resultant paths, spans the cluster.

When new nodes are added to a cluster, there is no need to update clients to point to the new nodes. The existence of the new nodes is transparent to the clients.

If you have a two-node cluster, you must configure cluster high availability (HA). For more information, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

You can create a cluster on a standalone node, called a single-node cluster. This configuration does not require a cluster network, and enables you to use the cluster ports to serve data traffic.

The nodes in a cluster communicate over a dedicated, physically isolated and secure Ethernet network. The cluster logical interfaces (LIFs) on each node in the cluster must be on the same subnet. For information about network management for cluster and nodes, see the *Clustered Data ONTAP Network Management Guide*.

What SVMs are

Storage Virtual Machines (SVMs, formerly known as Vservers) contain data volumes and one or more LIFs through which they serve data to the clients. Starting with clustered Data ONTAP 8.1.1, SVMs can either contain one or more FlexVol volumes, or a single Infinite Volume.

SVMs securely isolate the shared virtualized data storage and network, and each SVM appears as a single dedicated server to the clients. Each SVM has a separate administrator authentication domain and can be managed independently by its SVM administrator.

In a cluster, SVMs facilitate data access. A cluster must have at least one SVM to serve data. SVMs use the storage and network resources of the cluster. However, the volumes and LIFs are exclusive to the SVM. Multiple SVMs can coexist in a single cluster without being bound to any node in a cluster. However, they are bound to the physical cluster on which they exist.

A cluster can have one or more SVMs with FlexVol volumes and SVMs with Infinite Volume.

SVM with FlexVol volumes

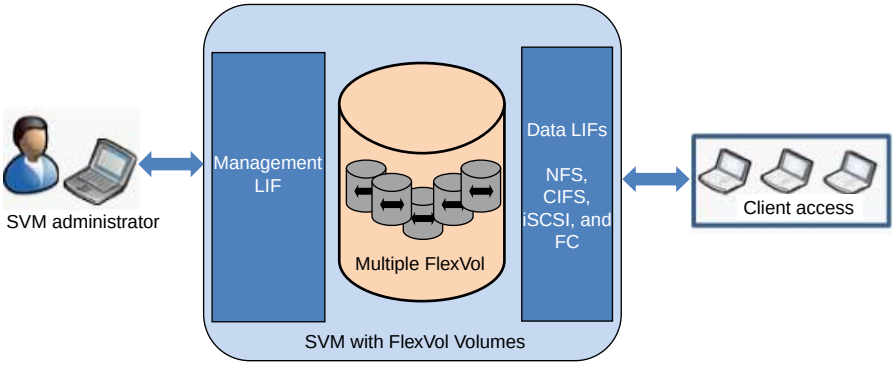

Each SVM with FlexVol volumes in a NAS environment presents a single directory hierarchical view and has a unique namespace. The namespace enables NAS clients to access data without specifying the physical location of the data. The namespace also enables the cluster and SVM administrators to manage distributed data storage as a single directory with multiple levels of hierarchy.

The volumes within each NAS SVM are related to each other through junctions and are mounted on junction paths. These junctions present the file system in each volume. The root volume of the SVM is a FlexVol volume that resides at the top level of the namespace hierarchy; additional volumes are mounted to the SVM root volume to extend the namespace. As volumes are created for the SVM, the root volume of the SVM contains junction paths.

SVMs with FlexVol volumes can contain files and LUNs. They provide file-level data access by using NFS and CIFS protocols for the NAS clients, and block-level data access by using iSCSI and Fibre Channel (FC) (FCoE included) for SAN hosts.

SVM with Infinite Volume

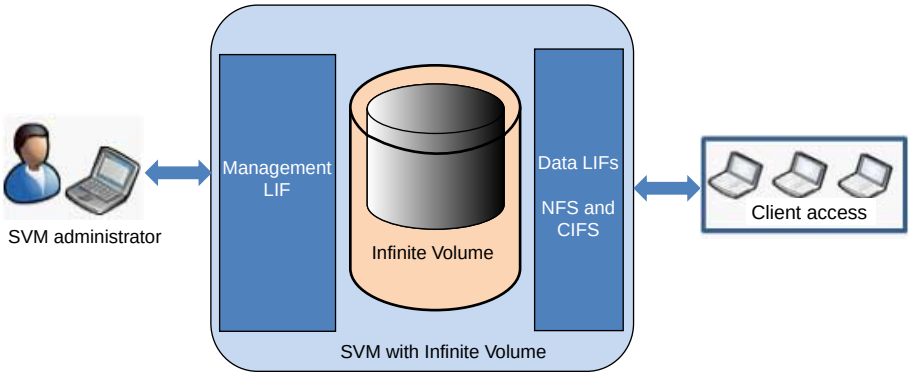

SVMs with Infinite Volume can contain only one Infinite Volume to serve data. Each SVM with Infinite Volume includes only one junction path, which has a default value of `/NS`. The junction provides a single mount point for the large namespace provided by the SVM with Infinite Volume. You cannot add more junctions to an SVM with Infinite Volume. However, you can increase the size of the Infinite Volume.

SVMs with Infinite Volume can contain only files. They provide file-level data access by using NFS and CIFS (SMB 1.0) protocols. SVMs with Infinite Volume cannot contain LUNs and do not provide block-level data access.

Note: The Data ONTAP command-line interface (CLI) continues to use the term *Vserver* in the output, and `vserver` as a command or parameter name has not changed.

Why you use SVMs

Storage Virtual Machines (SVMs, formerly known as Vservers) provide data access to clients without regard to physical storage or controller, similar to any storage system. When you use SVMs, they provide benefits such as nondisruptive operation, scalability, security and support unified storage.

SVMs have the following benefits:

- Nondisruptive operation
  SVMs can operate continuously and nondisruptively for as long as they are needed. SVMs help clusters to operate continuously during software and hardware upgrades, addition and removal of nodes, and all administrative operations.
- Scalability
  SVMs meet on-demand data throughput and the other storage requirements.
- Security
  Each SVM appears as a single independent server, which enables multiple SVMs to coexist while ensuring no data flows among them.
- Unified storage

SVMs can serve data concurrently through multiple data access protocols. SVMs provide file-level data access by using NAS protocols, such as CIFS and NFS, and block-level data access by using SAN protocols, such as iSCSI and FC (FCoE included). SVMs can serve data to SAN and NAS clients independently at the same time.

Note: SVMs with Infinite Volume can serve data only through NFS and CIFS (SMB 1.0) protocols.

- Delegation of management

  Each SVM can have its own user and administration authentication. SVM administrators can manage the SVMs that they are authorized to access. However, SVM administrators have privileges assigned by the cluster administrators.
- Easy management of large datasets

  With SVMs with Infinite Volume, management of large and unstructured data is easier as the SVM administrator has to manage one data container instead of many.

What logical storage is

Logical storage refers to the storage resources provided by Data ONTAP that are not tied to a physical resource.

Logical storage resources are associated with a Storage Virtual Machine (SVM, formerly known as Vserver), and they exist independently of any specific physical storage resource such as a disk, array LUN, or aggregate. Logical storage resources include volumes of all types and qtrees, as well as the capabilities and configurations you can use with these resources, such as Snapshot copies, deduplication, compression, and quotas.

For more information about SVMs, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators* and the *Clustered Data ONTAP System Administration Guide for SVM Administrators*.

How volumes work

Volumes are data containers that enable you to partition and manage your data. Understanding the types of volumes and their associated capabilities enables you to design your storage architecture for maximum storage efficiency and ease of administration.

Volumes are the highest-level logical storage object. Unlike aggregates, which are composed of physical storage resources, volumes are completely logical objects.

Data ONTAP provides two types of volumes: FlexVol volumes and Infinite Volumes. There are also volume variations, such as FlexClone volumes, FlexCache volumes, SnapLock volumes, data protection mirrors, and load-sharing mirrors. Not all volume variations are supported for both types of volumes. Data ONTAP efficiency capabilities, compression and deduplication, are supported for both types of volumes.

Volumes contain file systems in a NAS environment, and LUNs in a SAN environment.

Volumes are associated with one Storage Virtual Machine (SVM). The SVM is a virtual management entity, or server, that consolidates various cluster resources into a single manageable unit. When you

create a volume, you specify the SVM it is associated with. The type of the volume (FlexVol volume or Infinite Volume) is determined by an immutable SVM attribute.

Volumes have a language. The language of the volume determines the character set Data ONTAP uses to display file names and data for that volume. The default value for the language of the volume is the language of the SVM.

Volumes depend on their associated aggregates for their physical storage; they are not directly associated with any concrete storage objects, such as disks or RAID groups. If the cluster administrator has assigned specific aggregates to an SVM, then only those aggregates can be used to provide storage to the volumes associated with that SVM. This impacts volume creation, and also copying and moving FlexVol volumes between aggregates.

For more information about Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

For more information about SVMs, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

For more information about data protection mirrors, see the *Clustered Data ONTAP Data Protection Guide*.

For more information about physical storage resources such as aggregates, disks, and RAID groups, see the *Clustered Data ONTAP Physical Storage Management Guide*.

What LIFs are

A LIF (logical interface) is an IP address with associated characteristics, such as a role, a home port, a home node, a routing group, a list of ports to fail over to, and a firewall policy. You can configure LIFs on ports over which the cluster sends and receives communications over the network.

LIFs can be hosted on the following ports:

- Physical ports that are not part of interface groups
- Interface groups
- VLANs
- Physical ports or interface groups that host VLANs

While configuring SAN protocols such as FC on a LIF, it will be associated with a WWPN.

For more information about configuring WWPN to LIFs while using the FC protocol, see the *Clustered Data ONTAP SAN Administration Guide*.

The following figure illustrates the port hierarchy in a clustered Data ONTAP system:

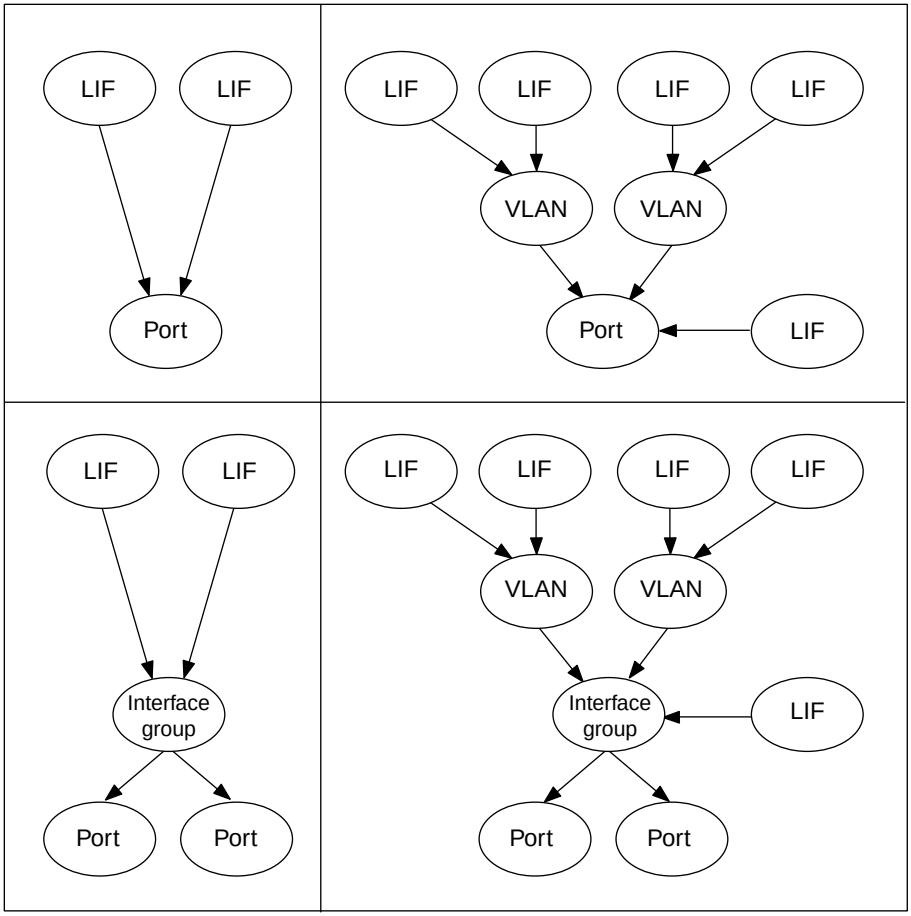

Types of network ports

Ports are either physical ports (NICs), or virtualized ports such as interface groups or VLANs. Interface groups treat several physical ports as a single port, while VLANs subdivide a physical port into multiple separate virtual ports.

Interface group A port aggregate containing two or more physical ports that act as a single trunk port. An interface group can be single-mode, multimode, or dynamic multimode.

VLAN A virtual port that receives and sends VLAN-tagged (IEEE 802.1Q standard) traffic. VLAN port characteristics include the VLAN ID for the port. The underlying physical port or interface group ports are considered VLAN trunk ports, and the connected switch ports must be configured to trunk the VLAN IDs.

> **Note:** The underlying physical port or interface group ports for a VLAN port can continue to host LIFs, which transmit and receive untagged traffic.

What the Snapshot copy reserve is

The Snapshot copy reserve sets a specific percent of the disk space for Snapshot copies. By default, the Snapshot copy reserve is 5 percent of the disk space for a FlexVol volume and 0 percent for aggregates.

The active file system cannot consume the Snapshot copy reserve space, but the Snapshot copy reserve, if exhausted, can use space in the active file system.

What an HA pair is

An HA pair is two storage systems (nodes) whose controllers are connected to each other directly. In this configuration, one node can take over its partner's storage to provide continued data service if the partner goes down.

You can configure the HA pair so that each node in the pair shares access to a common set of storage, subnets, and tape drives, or each node can own its own distinct set of storage.

The controllers are connected to each other through an HA interconnect. This allows one node to serve data that resides on the disks of its failed partner node. Each node continually monitors its partner, mirroring the data for each other’s nonvolatile memory (NVRAM or NVMEM). The interconnect is internal and requires no external cabling if both controllers are in the same chassis.

Takeover is the process in which a node takes over the storage of its partner. *Giveback* is the process in which that storage is returned to the partner. Both processes can be initiated manually or configured for automatic initiation.

How HA pairs relate to the cluster

HA pairs are components of the cluster, and both nodes in the HA pair are connected to other nodes in the cluster through the data and cluster networks. But only the nodes in the HA pair can takeover each other's storage.

Although the controllers in an HA pair are connected to other controllers in the cluster through the cluster network, the HA interconnect and disk-shelf connections are found only between the node and its partner and their disk shelves or array LUNs.

The HA interconnect and each node's connections to the partner's storage provide physical support for high-availability functionality. The high-availability storage failover capability does not extend to other nodes in the cluster.

> **Note:** Network failover does not rely on the HA interconnect and allows data network interfaces to failover to different nodes in the cluster outside the HA pair. Network failover is different than storage failover since it enables network resiliency across all nodes in the cluster.

Non-HA (or stand-alone) nodes are not supported in a cluster containing two or more nodes. Although single node clusters are supported, joining two separate single node clusters to create one cluster is not supported, unless you wipe clean one of the single node clusters and join it to the other

to create a two-node cluster that consists of an HA pair. For information on single node clusters, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

The following diagram shows two HA pairs. The multipath HA storage connections between the nodes and their storage are shown for each HA pair. For simplicity, only the primary connections to the data and cluster networks are shown.

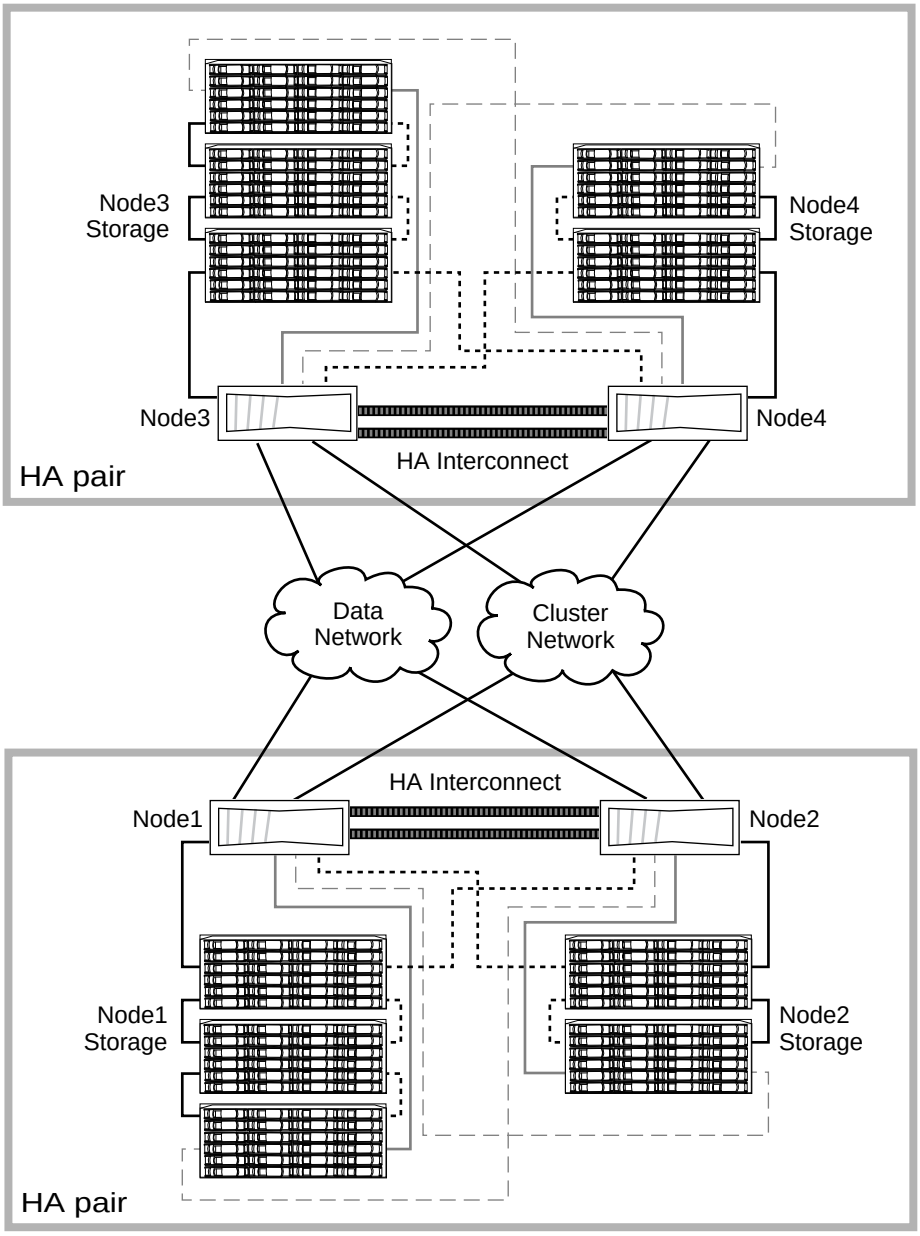

Key to storage connections

- Primary connection
- Redundant primary connection
- Standby connection
- Redundant standby connection

Possible storage failover scenarios in this cluster are as follows:

- Node1 fails and Node2 takes over Node1's storage.
- Node2 fails and Node1 takes over Node2's storage.
- Node3 fails and Node4 takes over Node3's storage.
- Node4 fails and Node3 takes over Node4's storage.

If Node1 *and* Node2 both fail, the storage owned by Node1 and Node2 becomes unavailable to the data network. Although Node3 and Node4 are clustered with Node1 and Node2, they do not have direct connections to Node1 and Node2's storage and cannot take over their storage.

When takeovers occur

Takeovers can be initiated manually or occur automatically when a failover event happens, depending on how you configure the HA pair. In some cases, takeovers occur automatically, regardless of configuration.

Takeovers can occur under the following conditions:

- A takeover is manually initiated.
- A node is in an HA pair with the default configuration for immediate takeover on panic, and that node undergoes a software or system failure that leads to a panic.

  By default, the node automatically performs a giveback, returning the partner to normal operation after the partner has recovered from the panic and booted up.
- A node that is in an HA pair undergoes a system failure (for example, a loss of power) and cannot reboot.

  Note: If the storage for a node also loses power at the same time, a standard takeover is not possible.
- A node does not receive heartbeat messages from its partner.

  This could happen if the partner experienced a hardware or software failure that did not result in a panic but still prevented it from functioning correctly.
- Hardware-assisted takeover is enabled and triggers a takeover when the remote management device (RLM or Service Processor) detects failure of the partner node.

Connections and components of an HA pair

Each node in an HA pair requires a network connection, an HA interconnect between the controllers, and connections both to its own disk shelves as well as its partner node's shelves.

Standard HA pair

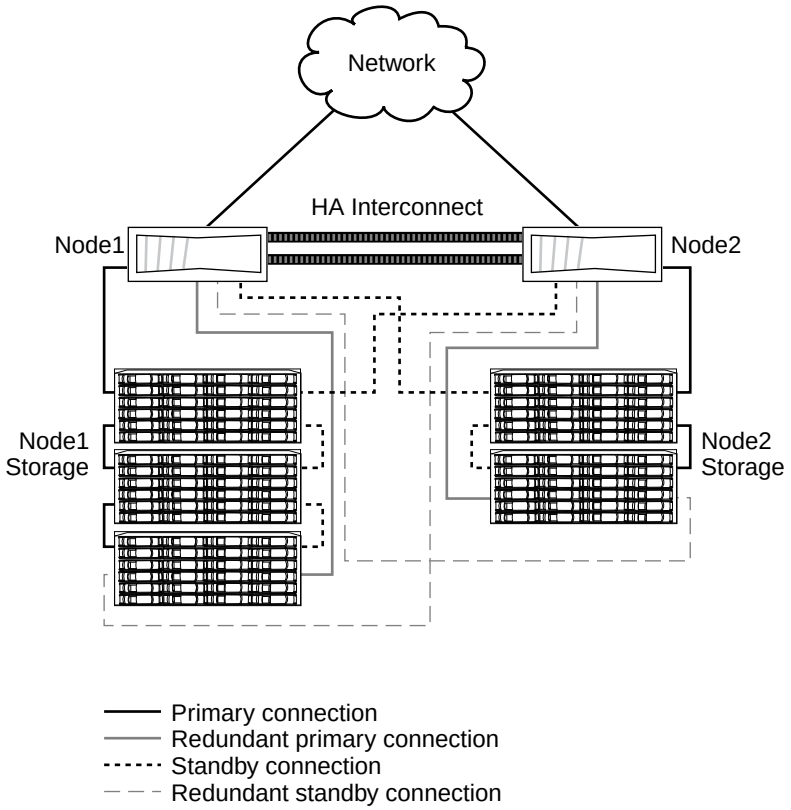

This diagram shows a standard HA pair with native disk shelves and multipath HA.

This diagram shows DS4243 disk shelves. For more information about cabling SAS disk shelves, see the *Universal SAS and ACP Cabling Guide* on the NetApp Support Site.

Understanding takeover and giveback

Takeover and giveback are the operations that let you take advantage of the HA configuration to perform nondisruptive operations and avoid service interruptions. Takeover is the process in which a

node takes over the storage of its partner. Giveback is the process in which the storage is returned to the partner.

About CIFS and SMB

Data ONTAP supports all of the most common file protocols, including the CIFS protocol to enable file sharing from host storage systems. When your system is first installed and CIFS is configured in Workgroup mode, a login named "administrator" is automatically created. You can use this login to access shares with a blank password.

The CIFS protocol is used to share files. CIFS is the method of transport for Windows Shares. CIFS is an extension of the Server Message Block (SMB) protocol, which is a file-sharing protocol used on Windows and UNIX systems. SMB runs over several different types of networks, including TCP/IP. For most purposes, SMB is superseded by CIFS.

Related concepts

CIFS license on page 143
Supported SMB clients and domain controllers on page 143

CIFS license

Your storage system requires a software license to enable CIFS service. This license is installed on the storage system at the factory per your order; therefore, you do not typically need to enter license codes when you initially configure your system. If CIFS license is not installed on the storage system, then System Manager does not list 'cifs' in the Licenses window (**Configuration > System Tools > Licenses**).

You need to enter license codes only if any of the following conditions applies:

- You purchased a storage system with a software version earlier than Data ONTAP 4.0 and you are upgrading it.
- You want to enable CIFS, which was not previously licensed for your storage system.
- You reinstalled your file system on an existing system that was not shipped with it installed.

In these cases, you are provided with the appropriate license codes when the software upgrade kit is shipped to you or when you are given instructions for obtaining the software upgrade over the Internet.

Supported SMB clients and domain controllers

Before you can use SMB with your Storage Virtual Machine (SVM), you need to know which SMB clients and domain controllers Data ONTAP supports.

For the latest information about which SMB clients and domain controllers Data ONTAP supports, see the Interoperability Matrix at *support.netapp.com/matrix*.

What a Snapshot copy is

A Snapshot copy is a frozen, read-only image of a flexible volume, or an aggregate that captures the state of the file system at a point in time. Snapshot copies are your first line of defense to backup and restore data.

When Snapshot copies are created, Data ONTAP maintains a configurable Snapshot copy schedule that creates and deletes Snapshot copies automatically for each volume. You can also create and delete Snapshot copies manually.

You can store up to 255 Snapshot copies at one time on each volume.

You can specify the percentage of disk space that Snapshot copies can occupy. The default space reserved for Snapshot copies is zero percent for SAN and VMware volumes. For NAS volumes, it is five percent on storage systems running Data ONTAP 8.1.

Description of cluster object windows and dialog boxes

You can view all your clusters and cluster objects from the respective storage object page. You can also view the details from the corresponding storage object details page.

Clusters page

The Clusters page enables you to add clusters and to view detailed information about the clusters that you are monitoring.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

- *Command buttons* on page 144
- *Clusters list* on page 145
- *Filters pane* on page 146

Command buttons

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

Add	Opens the Add Cluster dialog box, which enables you to add clusters.
Edit	Opens the Edit Cluster dialog box, which enables you to edit the settings of the selected cluster.
Remove	Removes the selected cluster and all the associated events and storage objects. After the cluster is removed, it is no longer monitored.
View Monitoring Status	Enables you to view the monitoring statuses of the selected clusters by navigating to the Data Sources page.

Refresh List	Refreshes the clusters list and the properties associated with the cluster.
Export	Enables you to export the details of all the monitored clusters to a comma-separated values (`.csv`) file.

Clusters list

The Clusters list displays, in tabular format, the properties of all the discovered clusters. You can use the column filters to customize the data that is displayed:

Status	Displays the current status of the cluster. The status can be Critical ( ), Error ( ), Warning ( ), or Normal ( ). You can move the pointer over the status to view more information about the event or events generated for the cluster. If the status of the cluster is based on a single event, you can view information such as the event name, time and date when the event was generated, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event. If the status of the cluster is based on multiple events of the same severity, the top three events are displayed along with information such as the event name, time and date when the events are generated, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.
Cluster	Displays the name of the cluster.
Communication Status	Displays whether the cluster is reachable or not. The communication status is displayed as Good if the cluster is reachable. If the cluster is not reachable or if the login credentials are invalid, the communication status is displayed as Not Reachable.
System Health	Displays high-level information about the status of the cluster, which is calculated based on the status of the various cluster subsystems. Possible values are: OK, OK with suppressed, Degraded, and Components not reachable. These values are determined by the health monitors in Data ONTAP.
Host Name or IP Address	Displays the FQDN, short name, or the IP address of the cluster-management LIF that is used to connect to the cluster.
FQDN	Displays the Fully Qualified Domain Name (FQDN) of the cluster. By default, this column is hidden.

OS Version	Displays the Data ONTAP version that the cluster is running. If the nodes in the cluster are running different versions of Data ONTAP, then the earliest Data ONTAP version is displayed.
Node Count	Displays the number of nodes that belong to the cluster. This column is nonsortable.
Last Refreshed Time	Displays the timestamp when the monitoring samples of the cluster were last collected. This column is nonsortable.
Serial Number	Displays the serial number of the cluster.
Contact	Displays the contact information of the cluster. By default, this column is hidden.
Location	Displays the location of the cluster.

Filters pane

The Filters pane enables you to set filters to customize the display of information in the clusters list. You can select filters in the Status, Communication Status, System Health, and Annotation columns.

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

Related tasks

Adding clusters on page 120
Editing clusters on page 122
Removing clusters on page 123
Exporting data to CSV files on page 18

Cluster details page

You can use the Cluster details page to view detailed information about the selected cluster that is monitored by Unified Manager, such as the health, capacity, and configuration details. You can also view information about the logical interfaces (LIFs), nodes, disks, related devices, and related alerts for that cluster.

- *Command buttons* on page 147
- *Health tab* on page 147
- *Capacity tab* on page 148
- *Configuration tab* on page 149
- *LIFs tab* on page 151
- *Nodes tab* on page 151
- *Disks tab* on page 153
- *Related Devices pane* on page 154

- *Related Alerts pane* on page 155
- *Annotations pane* on page 155

Command buttons

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

Actions

- Add Alert
  Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.
- Edit Cluster
  Opens the Edit Cluster dialog box, which enables you to edit the settings of the selected cluster.
- Rediscover
  Initiates a manual refresh of a cluster, which enables Unified Manager to discover recent changes to the cluster. If Unified Manager is paired with OnCommand Workflow Automation, this operation also reacquires WFA's cached data.
  After the rediscovery operation is initiated, a link to the associated job details is displayed to enable tracking the job status.

View Clusters Enables you to navigate to the Clusters page.

Health tab

The Health tab displays detailed information about data availability and data capacity issues of various cluster objects such as nodes, Storage Virtual Machines (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 the filtered list of the objects. For example, you can click the SVM capacity graph that displays warnings to view the filtered list of SVMs. The list contains SVMs that have volumes or qtrees that have capacity issues with severity as warning. You can also click the SVMs availability graph that displays warnings to view the list of SVMs that have availability issues with severity as warning.

Availability 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 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 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 cluster. For example, information is displayed about aggregates that are likely to breach the set threshold values.

Capacity tab

The Capacity tab displays detailed information about the capacity of the selected cluster:

Capacity The Capacity area displays details about the used 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.
- 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.
- SSD Tier
  Displays the total space of the solid-state disks (SSDs) that are added to the cluster.

Capacity Breakout by Disk Type The Capacity Breakout by Disk Type area displays detailed information about the disk capacity of 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 distribution of available, used, and spare capacity disks. The dotted line represents the total size of the spare disk capacity available for a disk type.
- Unassigned Disks

Displays the number of unassigned disks in the cluster.

- Cache

  Displays the total size of spare disk capacity.

Aggregates with Capacity Issues list

The Aggregates with Capacity Issues list displays, in tabular format, details about the used 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 severities of Error and Critical, only the Critical severity is displayed.

- Aggregate

  Displays the name of the aggregate.

- Used Data Capacity

  Displays, as a graph, 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

The Configuration tab displays details about the selected cluster, such as the IP address, serial number, contact, and location information of the cluster:

Overview

- Management LIF

Displays the cluster-management LIF that Unified Manager uses to connect to the cluster. The operational status of the LIF is also displayed.

- Host Name or IP Address

  Displays the FQDN, short name, or the IP address of the cluster-management LIF that Unified Manager uses to connect to the cluster.
- FQDN

  Displays the Fully Qualified Domain Name (FQDN) of the cluster.
- OS Version

  Displays the Data ONTAP version that the cluster is running. If the nodes in the cluster are running different versions of Data ONTAP, then the earliest Data ONTAP version is displayed.
- Serial Number

  Displays the serial number of the cluster.
- Contact

  Displays the contact information of the cluster.
- Location

  Displays the location of the cluster.

Nodes

- Availability

  Displays the number of nodes that are up ( ) or down ( ) in the cluster.
- OS Versions

  Displays the Data ONTAP versions that the nodes are running. Also, displays the number of nodes running a particular version of Data ONTAP. For example, 8.2 (2), 8.1 (1) specifies that two nodes are running Data ONTAP 8.2 and one node is running Data ONTAP 8.1.

Storage Virtual Machines

- Availability

  Displays the number of SVMs that are up ( ) or down ( ) in the cluster.

LIFs

- Availability

  Displays the number of non-data LIFs that are up ( ) or down ( ) in the cluster.
- Cluster-Management LIFs

  Displays the number of cluster-management LIFs.
- Node-Management LIFs

  Displays the number of node-management LIFs.
- Cluster LIFs

  Displays the number of cluster LIFs.
- Intercluster LIFs

  Displays the number of intercluster LIFs.

Protocols	• Data Protocols Displays the list of licensed data protocols that are enabled for the cluster. The data protocols include iSCSI, CIFS, NFS, and FC and FCoE.

LIFs tab

The 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.
Failover Group	Displays the name of the failover group.

Nodes tab

The 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 environments, 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.
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 port details. 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
  The firmware revision of the FC/FCoE port.
- Status
  Displays the current state of the port. The possible states are Online, Offline, or Unknown ( ).

You can view the port-related events form the Events list. You can also view the associated LIF details, such as LIF name, operational status, the IP address or the WWPN, protocols, name of the SVM associated with the LIF, current port, failover policy and failover group.

Disks tab

The 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, and Array LUN), and the state of the disks. You can also view other details, such as the number of broken disks, spare disks, and unassigned 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—Broken, Aggregate, Spare, Unknown, or Unassigned. By default, this column is sorted to display the states in the following order: Broken, Spare, Aggregate, Unknown, and Unassigned.

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 Aggregate	Displays the name of aggregate which is impacted due to the failed disk. By clicking the aggregate name, you can view the aggregate details in the Aggregate details page. If there is no failed disk, no value is displayed in this column.
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, or FCAL.
Effective Type	Displays the disk type assigned by Data ONTAP. Certain Data ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and spare management. Data ONTAP assigns an effective disk type for each disk type.
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. By default, this column is hidden.
Bay	Displays the ID of the bay where the disk is located. By default, this column is hidden.

Related Devices pane

The Related Devices pane enables you to view and navigate to the nodes, SVMs, and aggregates that are related to the cluster:

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 the 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 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.

Annotations pane

The Annotations pane enables you to view the annotation for the selected cluster.

Related tasks

Viewing the cluster list and details on page 124
Adding an alert on page 77
Editing clusters on page 122

Add Cluster dialog box

You can add an existing cluster to monitor the cluster and obtain information such as its health, capacity, and configuration.

You can add a cluster by specifying the following options:

Host Name or IP Address	Enables you to specify the FQDN, short name, or the IP address of the cluster-management LIF that is used to connect to the cluster.
User Name	Enables you to specify a user name that can be used to log in to the cluster.
Password	Enables you to specify a password for the specified user name.
Protocol	Enables you to specify the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.
Port	Enables you to specify the port number of the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Command buttons

The command buttons enable you to perform the following tasks:

Add	Adds an existing cluster and closes the Add Cluster dialog box.
Cancel	Discards the changes and closes the Add Cluster dialog box.

Related tasks

Adding clusters on page 120

Edit Cluster dialog box

The Edit Cluster dialog box enables you to modify the settings of an existing cluster, including the IP address, port, and protocol. For example, you can change the protocol from HTTP to HTTPS.

You can edit the following fields:

Host Name or IP Address	Enables you to specify the FQDN, short name, or the IP address of the cluster-management LIF that is used to connect to the cluster.
User Name	Enables you to specify a user name that can be used to log in to the cluster.
Password	Enables you to specify a password for the specified user name.
Protocol	Enables you to specify the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.
Port	Enables you to specify the port number of the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Command buttons

The command buttons enable you to perform the following tasks:

Save	Saves the changes made to the settings of an existing cluster and closes the Edit Cluster dialog box.
Cancel	Discards the changes and closes the Edit Cluster dialog box.

Related tasks

Editing clusters on page 122

Nodes page

The Nodes page enables you to view detailed information about the nodes in a selected cluster.

- *Command button* on page 156
- *Nodes list* on page 157
- *Filters pane* on page 158

Command button

The command button enables you to perform the following task for a selected node:

Export Enables you to export the details of all the monitored nodes to a comma-separated values (`.csv`) file.

Nodes list

The Nodes list displays, in tabular format, the properties of all the discovered nodes in a cluster. You can use the column filters to customize the data that is displayed:

Status	Displays the current status of the node. The status can be Critical ( ), Error ( ), Warning ( ), or Normal ( ).
Node	Displays the name of the node.
State	Displays the state of the node. The state can be Up or Down. By default, this column is hidden.
HA State	Displays the state of the HA pair. The state can be Error, Warning, Normal, or Not applicable.
Down Time	Displays the time that has elapsed or the timestamp since the node is offline. If the time elapsed exceeds a week, the timestamp when the node went offline is displayed. By default, this column is hidden.
Cluster	Displays the name of the cluster to which the node belongs.
Model	Displays the model of the node.
OS version	Displays the Data ONTAP version that the node is running.
Owner	Displays the owner name of the node.
Serial Number	Displays the serial number of the node. By default, this column is hidden.
Firmware Version	Displays the firmware version number of the node. By default, this column is hidden.
Contact	Displays the contact information of the node. By default, this column is hidden.
Location	Displays the location of the node.
Aggregate Used Capacity	Displays the amount of space (in GB or MB) used for data in the node's aggregates.
Aggregate Total Capacity	Displays the total space (in GB or MB) available for data in the node's aggregates.
Usable Raw Capacity	Displays the amount of space (in GB or MB) that is usable in the node.
Usable Spare Capacity	Displays the amount of available space (in GB or MB) in the node to enhance the aggregate capacity.

Total Raw Capacity	Displays the capacity of every unformatted disk in the node before right-sizing and RAID configuration.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the nodes list. You can select filters related to the Status, State, and HA State columns.

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

Related tasks

Exporting data to CSV files on page 18

Storage Virtual Machines page

The Storage Virtual Machines page enables you to view detailed information about the Storage Virtual Machines (SVMs) that you are monitoring.

- *Command button* on page 158
- *SVMs list* on page 158
- *Filters pane* on page 159

Command button

The command button enables you to perform the following task for a selected SVM:

Export	Enables you to export the details of all the monitored SVMs to a comma-separated values (`.csv`) file.

SVMs list

The SVMs list displays, in tabular format, the properties of all the discovered SVMs. You can use the column filters to customize the data that is displayed:

Status	Displays the current status of the SVM. The status can be Critical ( ), Error ( ), Warning ( ), or Normal ( ). You can move the pointer over the status to view more information about the event or events generated for the SVM. If the status of the SVM is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the View Details button to view more information about the event.

	If the status of the SVM is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the View All Events link to view the list of generated events.
Storage Virtual Machine	Displays the name of the SVM. You can move the pointer over each SVM to view information such as the last generated event, cluster to which the SVM belongs, volume type of the SVM, allowed protocols, and space allocated in the SVM. You can also view the details of related objects such as the cluster to which the SVM belongs, all the SVMs that belong to the cluster, and the volumes that belong to the SVM.
State	Displays the current state of the SVM. The state can be Running, Stopped, Starting, or Stopping.
Cluster	Displays the name of the cluster to which the SVM belongs.
Allowed Volume Type	Displays the type of volume that can be created in the SVM. The type can be InfiniteVol or FlexVol.
Available Data Capacity	Displays the available data capacity (in MB, GB, and so on) of all the volumes in the SVM.
Total Data Capacity	Displays the total data capacity (in MB, GB, and so on) of all the volumes in the SVM.
Root Volume	Displays the name of the root volume of the SVM. By default, this column is hidden.
NIS State	Displays the state of the Network Information Service (NIS). The state can be Enabled, Disabled, or Not Configured. By default, this column is hidden.
NIS Domain	Displays the NIS domain name. This column is blank when the NIS server is disabled or is not configured. By default, this column is hidden.
DNS State	Displays the state of the Domain Name System (DNS). The state can be Enabled, Disabled, or Not Configured. By default, this column is hidden.
DNS Domain	Displays the DNS domain name. By default, this column is hidden.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the SVMs list. You can select filters related to the Status, State, and Annotation columns.

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

Related tasks

Viewing the details of SVMs with Infinite Volume on page 294
Exporting data to CSV files on page 18

Storage Virtual Machine details page

You can use the Storage Virtual Machine details page to view detailed information about the selected Storage Virtual Machine (SVM, formerly known as Vserver) that is monitored by Unified Manager, such as its health, capacity, configuration, data policies, logical interfaces (LIFs), LUNs, qtrees, and user and user group quotas. You can also view information about the related objects and related alerts for the SVM.

Note: You can monitor only data SVMs.

- *Command buttons* on page 160
- *Health tab* on page 161
- *Capacity tab* on page 161
- *Configuration tab* on page 164
- *LIFs tab* on page 165
- *Qtrees tab* on page 166
- *User and Group Quotas tab* on page 167
- *NFS Exports tab* on page 169
- *CIFS Shares tab* on page 170
- *SAN tab* on page 172
- *Data Policy tab* on page 172
- *Related Devices pane* on page 174
- *Related Alerts pane* on page 174
- *Annotations pane* on page 174

Command buttons

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

Actions	• Add Alert Enables you to add an alert to the selected SVM. • Edit Thresholds Enables you to edit the SVM thresholds. **Note:** This button is enabled only for SVM with Infinite Volume.
View Storage Virtual Machines	Enables you to navigate to the Storage Virtual Machines page.

Health tab

The Health tab displays detailed information about data availability and data capacity issues of various objects such as volumes, aggregates, NAS LIFs, SAN LIFs, LUNs, protocols, services, NFS exports, and CIFS shares. Availability issues are related to the data-serving capability of the SVM objects. Capacity issues are related to the data-storing capability of the SVM objects.

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 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 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.

Capacity tab

The Capacity tab displays detailed information about the data capacity of the selected SVM.

The following information is displayed for an SVM with FlexVol volume:

Capacity The Capacity area displays details about the used and available capacity allocated from all volumes:

- Total Capacity
  Displays the total capacity (in MB, GB, and so on) of the SVM.
- Used
  Displays the space used by data in the volumes that belong to the SVM.

- Guaranteed Available

  Displays the guaranteed available space for data that is available for volumes in the SVM.
- Unguaranteed

  Displays the available space remaining for data that is allocated for thinly provisioned volumes in the SVM.

Volumes with Capacity Issues

The Volumes with Capacity Issues list displays, in tabular format, details about the volumes that have capacity issues:

- Status

  Indicates that the volume has a capacity-related issue of a certain severity.

  You can move the pointer over the status to view more information about the capacity-related event or events generated for the volume.

  If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use the **View Details** button to view more information about the event.

  If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

  Note: A volume can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a volume has two events with severities of Error and Warning, only the Error severity is displayed.
- Volume

  Displays the name of the volume.
- Used Data Capacity

  Displays, as a graph, information about the volume capacity usage (in percentage).
- Days to Full

  Displays the estimated number of days remaining before the volume reaches full capacity.
- Thin Provisioned

  Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.
- Aggregate

  Displays the name of the aggregate that contains the volume.

The following information is displayed for an SVM with Infinite volume:

Capacity

Displays the following capacity-related details:

- Percentage of used and free data capacity
- Percentage of used and free Snapshot capacity
- Snapshot Overflow
  Displays the data space that is consumed by the Snapshot copies.
- Used
  Displays the space used by data in the SVM with Infinite Volume.
- Warning
  Indicates that the space in the SVM with Infinite Volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.
- Error
  Indicates that the space in the SVM with Infinite Volume if full. If this threshold is breached, the Space Full event is generated.

Other Details

- Total Capacity
  Displays the total capacity in the SVM with Infinite Volume.
- Data Capacity
  Displays used data capacity, available data capacity, and Snapshot overflow capacity details of the SVM with Infinite Volume.
- Snapshot Reserve
  Displays the used and free details of the Snapshot reserve.
- System Capacity
  Displays the used system capacity and available system capacity in the SVM with Infinite Volume.
- Thresholds
  Displays the nearly full and full thresholds of the SVM with Infinite Volume.

Storage Class Capacity Details

Displays information about the capacity usage in your storage classes. This information is displayed only if you have configured storage classes for your SVM with Infinite Volume.

Storage Virtual Machine Storage Class Thresholds

Displays the following thresholds (in percentage) of your storage classes:

- Nearly Full Threshold
  Specifies the percentage at which a storage class in an SVM with Infinite Volume is considered to be nearly full.
- Full Threshold
  Specifies the percentage at which the storage class in an SVM with Infinite Volume is considered full.
- Snapshot Usage Limit
  Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.

Configuration tab

The Configuration tab displays configuration details about the selected SVM, such as its cluster, root volume, the type of volumes it contains (Infinite Volume or FlexVol volumes), and the policies created on the SVM:

Overview

- Cluster
  Displays the name of the cluster to which the SVM belongs.
- Allowed Volume Type
  Displays the type of volumes that can be created in the SVM. The type can be InfiniteVol or FlexVol.
- Root Volume
  Displays the name of the root volume of the SVM.
- Allowed Protocols
  Displays the type of protocols that can be configured on the SVM. Also, indicates if a protocol is up ( ), down ( ), or is not configured ( ).

Data LIFs

- NAS
  Displays the number of NAS LIFs that are associated with the SVM. Also, indicates if the LIFs are up ( ) or down ( ).
- SAN
  Displays the number of SAN LIFs that are associated with the SVM. Also, indicates if the LIFs are up ( ) or down ( ).
- Junction Path
  Displays the path on which the Infinite Volume is mounted. Junction path is displayed for an SVM with Infinite Volume only.
- Storage Classes
  Displays the storage classes associated with the selected SVM with Infinite Volume. Storage classes are displayed for an SVM with Infinite Volume only.

Management LIFs

- Availability
  Displays the number of management LIFs that are associated with the SVM. Also, indicates if the management LIFs are up ( ) or down ( ).

Policies

- Snapshots
  Displays the name of the Snapshot policy that is created on the SVM.
- Export Policies
  Displays either the name of the export policy if a single policy is created or displays the number of export policies if multiple policies are created.
- Data Policy

Displays whether a data policy is configured for the selected SVM with Infinite Volume.

Services

- Type
  Displays the type of service that is configured on the SVM. The type can be Domain Name System (DNS) or Network Information Service (NIS).
- State
  Displays the state of the service, which can be Up ( ), Down ( ), or Not Configured ( ).
- Domain Name
  Displays the fully qualified domain names (FQDNs) of the DNS server for the DNS services or NIS server for the NIS services. When the NIS server is enabled, the active FQDN of the NIS server is displayed. When the NIS server is disabled, the list of all the FQDNs are displayed.
- IP Address
  Displays the IP addresses of the DNS or NIS server. When the NIS server is enabled, the active IP address of the NIS server is displayed. When the NIS server is disabled, the list of all the IP addresses are displayed.

LIFs tab

The LIFs tab displays details about the data LIFs that are created on the selected SVM:

LIF	Displays the name of the LIF that is created on the selected SVM.
Operational Status	Displays the operational status of the LIF, which can be Up ( ), Down ( ), or Unknown ( ). The operational status of a LIF is determined by the status of its physical ports.
Administrative Status	Displays the administrative status of the LIF, which can be Up ( ), Down ( ), or Unknown ( ). The administrative status of a LIF is controlled by the storage administrator to make changes to the configuration or for maintenance purposes. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.
IP Address / WWPN	Displays the IP address for Ethernet LIFs and the World Wide Port Name (WWPN) for FC LIFs.
Protocols	Displays the list of data protocols that are specified for the LIF, such as CIFS, NFS, iSCSI, FC/FCoE, and FlexCache. For Infinite Volume, the SAN protocols are not applicable.

Role	Displays the LIF role. The roles can be Data or Management.
Home Port	Displays the physical port to which the LIF was originally associated.
Port Set	Displays the port set to which the LIF is mapped.
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.
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.
Failover Group	Displays the name of the failover group.

Qtrees tab

The Qtrees tab displays details about qtrees and their quotas:

Note: The Qtrees tab is not displayed for an SVM with Infinite Volume.

Status

Displays the current status of the qtree. The status can be Critical (), Error (), Warning (), or Normal ().

You can move the pointer over the status icon to view more information about the event or events generated for the qtree.

If the status of the qtree is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the qtree is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

Note: A qtree can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a qtree has two events with severities of Error and Warning, only the Error severity is displayed.

Qtree

Displays the name of the qtree.

Volume	Displays the name of the volume that contains the qtree. You can move the pointer over the volume name to view more information about the volume.
Quota Set	Indicates whether a quota is enabled or disabled on the qtree.
Disk Used %	Displays the percentage of disk space used. If a disk hard limit is set, this value is based on the disk hard limit. If the quota is set without a disk hard limit, the value is based on the volume data space. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then “Not applicable” is displayed.
Disk Hard Limit	Displays the maximum amount of disk space allocated for the qtree. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.
Disk Soft Limit	Displays the amount of disk space allocated for the qtree before a warning event is generated. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.
Disk Threshold	Displays the threshold value set on the disk space. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk threshold limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.
Files Used %	Displays the percentage of files used in the qtree. If the file hard limit is set, this value is based on the file hard limit. No value is displayed if the quota is set without a file hard limit. The value is displayed as “Not applicable” if the quota is not set or if quotas are off on the volume to which qtree belongs.
File Hard Limit	Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.
File Soft Limit	Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

User and Group Quotas tab

Displays details about the user and user group quotas for the selected SVM. You can view information such as the status of the quota, name of the user or user group, soft and hard limits set on the disks and files, amount of disk space and number of files used, and the disk threshold value. You can also change the email address associated with a user or user group.

Edit Email Address

Opens the Edit Email Address dialog box, which displays the current email address of the selected user or user group. You can modify the email address. If the **Edit Email Address** field is blank, the default rule is used to generate an email address for the selected user or user group.

If more than one user has the same quota, the names of the users are displayed as comma-separated values. Also, the default rule is not used to generate the email address; therefore, you must provide the required email address for notifications to be sent.

Configure Email Rules

Enables you to create or modify rules to generate an email address for the user or user group quotas that are configured on the SVM. A notification is sent to the specified email address when there is a quota breach.

Status

Displays the current status of the quota. The status can be Critical (), Warning (), or Normal ().

You can move the pointer over the status icon to view more information about the event or events generated for the quota.

If the status of the quota is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the quota is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

> **Note:** A quota can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a quota has two events with severities of Error and Warning, only the Error severity is displayed.

User or Group

Displays the name of the user or user group. If more than one user has the same quota, the names of the users are displayed as comma-separated values.

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

Type

Specifies if the quota is for a user or a user group.

Volume or Qtree

Displays the name of the volume or qtree on which the user or user group quota is specified.

	You can move the pointer over the name of the volume or qtree to view more information about the volume or qtree.
Disk Used %	Displays the percentage of disk space used. The value is displayed as “Not applicable” if the quota is set without a disk hard limit.
Disk 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 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 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 the NFS exports such as its status, the path associated with the volume (Infinite Volumes or FlexVol volumes), access levels of clients to the NFS exports, and the export policy defined for the volumes that are exported. NFS Export will not be displayed for 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.

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.
Junction Path Active	Displays whether the path to access the mounted volume is active or inactive.

Volume / Storage Virtual Machine	Displays the name of the volume, if the volume being exported is a FlexVol volume. For Infinite Volumes, the name of the SVM containing the Infinite Volume is displayed.
Volume State	Displays the state of the volume that is being exported. The state can be Offline, Online, or Restricted. • Offline Read or write access to the volume is not allowed. • Online Read and write access to the volume is allowed. • Restricted Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
Security Style	Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed. • UNIX (NFS clients) Files and directories in the volume have UNIX permissions. • Unified Files and directories in the volume have a unified security style. • NTFS (CIFS clients) Files and directories in the volume have Windows NTFS permissions. • Mixed Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.
UNIX Permission	Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.
Export Policy	Displays the rules that define the access permission for volumes that are exported. You can click the link to view details about the rules associated with the export policy such as the authentication protocols and the access permission.

CIFS Shares tab

Displays information about the CIFS shares on the selected SVM. You can view information such as the status of the CIFS share, share name, path associated with the SVM, the status of the junction path of the share, containing object, state of the containing volume, security data of the share, and export policies defined for the share. You can also determine whether an equivalent NFS path for the CIFS share exists.

View User Mapping	Launches the User Mapping dialog box. You can view the details of user mapping for the SVM.

Show ACL

Launches the Access Control dialog box for the share.

You can view user and permission details for the selected share.

Status

Displays the current status of the share. The status can be Normal ( ) or Error ( ).

Share Name

Displays the name of the CIFS share.

Path

Displays the junction path on which the share is created.

Junction Path Active

Displays whether the path to access the share is active or inactive.

Containing Object

Displays the name of the containing object to which the share belongs. The containing object can be a volume or a qtree.

By clicking the link, you can view details about the containing object in the respective Details page. If the containing object is a qtree, links are displayed for both qtree and volume.

Volume State

Displays the state of the volume that is being exported. The state can be Offline, Online, or Restricted.

- Offline
  Read or write access to the volume is not allowed.
- Online
  Read and write access to the volume is allowed.
- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Security

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

- UNIX (NFS clients)
  Files and directories in the volume have UNIX permissions.
- NTFS (CIFS clients)
  Files and directories in the volume have Windows NTFS permissions.
- Mixed
  Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

Export Policy

Displays the name of the export policy applicable on the share. If an export policy is not specified for the SVM, the value is displayed as Not Enabled.

	You can click the link to view details about the rules associated with the export policy, such as access protocols and permissions. The link is disabled if the export policy is disabled for the selected SVM.
NFS Equivalent	Specifies whether there is an NFS equivalent for the share.

SAN tab

Displays details about LUNs, initiator groups, and initiators for the selected SVM. By default, the LUNs view is displayed. You can view details about the initiator groups in the Initiator Groups tab and details about initiators in the Initiators tab.

LUNs	Displays details about the LUNs that belong to the selected SVM. You can view information such as the LUN name, LUN state (online or offline), the name of the file system (volume or qtree) that contains the LUN, the type of host operating system, the total data capacity and serial number of the LUN. You can also view information whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group. You can also view the initiator groups and initiators that are mapped to the selected LUN.
Initiator Groups	Displays details about initiator groups. You can view details such as the name of the initiator group, the access paths to which the initiator group is connected (one, many, or no paths), the type of host operating system used by all of the initiators in the group, and the supported protocol. You can view if 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, the LIFs that are mapped through the target portal are not displayed. The total number of initiators and LUNs that are mapped to an initiator group is displayed. You can also view the LUNs and initiators that are mapped to the selected initiator group.
Initiators	Displays the name and type of the initiator and the total number of initiator groups mapped to this initiator for the selected SVM. You can also view the LUNs and initiator groups that are mapped to the selected initiator group.

Data Policy tab

The Data Policy tab enables you to create, modify, activate, or delete one or more rules in a data policy. You can also import the data policy into the Unified Manager database and export the data policy to your computer:

Note: The Data Policy tab is displayed only for SVMs with Infinite Volume.

Rules list

Displays the list of rules. By expanding the rule, you can view the corresponding matching criteria of the rule and the storage class where the content is placed based on the rule.

The default rule is the last rule in the list. You cannot change the order of the default rule.

- Matching Criteria

  Displays the conditions for the rule. For example, a rule can be “File path starts with `/eng/nightly`”.

  Note: The file path must always start with a junction path.

- Content Placement

  Displays the corresponding storage class for the rule.

Rule Filter

Enables you to filter rules associated with a specific storage class listed in the list.

Action buttons

- Create

  Opens the Create Rule dialog box, which enables you to create a new rule for your data policy.
- Edit

  Opens the Edit Rule dialog box, which enables you to modify rule properties such as directory paths, file types, and owners.
- Delete

  Deletes the selected rule.
- Move Up

  Moves the selected rule up in the list. However, you cannot move the default rule up in the list.
- Move Down

  Moves the selected rule down the list. However, you cannot move the default rule down the list.
- Activate

  Activates the rules and changes made to the data policy in the SVM with Infinite Volume.
- Reset

  Resets all changes made to the data policy configuration.
- Import

  Imports a data policy configuration from a file.
- Export

  Exports a data policy configuration to a file.

Related Devices area

The Related Devices area enables you to view and navigate to the LUNs, CIFS shares, and the user and user group quotas that are related to the qtree:

LUNs	Displays the total number of the LUNs associated with the selected qtree.
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 Devices pane

The Related Devices pane enables you to view the cluster, aggregates, and volumes that are related to the SVM:

Cluster	Displays the health status of the cluster to which the SVM belongs.
Aggregates	Displays the number of aggregates that belong to the selected SVM. The health status of the aggregates is also displayed, based on the highest severity level. For example, if an SVM contains ten aggregates, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.
Assigned Aggregates	Displays the number of aggregates that are assigned to an SVM. The health status of the aggregates is also displayed, based on the highest severity level.
Volumes	Displays the number and capacity of the volumes that belong to the selected SVM. The health status of the volumes is also displayed, based on the highest severity level.

Related 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.

Annotations pane

The Annotations pane enables you to view the annotation for the selected SVM.

Related tasks

Viewing the details of SVMs with Infinite Volume on page 294
Viewing the constituents of an Infinite Volume on page 294
Editing the Infinite Volume threshold settings on page 295
Creating rules on page 298

Viewing rules on page 301
Previewing changes to your data policy on page 304
Adding an alert on page 77

Related references

Export Policy Rules dialog box on page 204

Aggregates page

The Aggregates page displays information about the aggregates that are monitored and enables you to view and modify the threshold settings.

- *Command buttons* on page 175
- *Aggregates list* on page 175
- *Filters pane* on page 177

Command buttons

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

Edit Thresholds	Displays the Edit Aggregate Thresholds dialog box, which enables you to edit the threshold settings for one or more aggregates.
Export	Enables you to export the details of all the monitored aggregates to a comma-separated values (`.csv`) file.

Aggregates list

The Aggregates list displays, in tabular format, the properties of all the discovered aggregates. You can use the column filters to customize the data that is displayed:

Status

Displays the current status of the aggregate. The status can be Critical ( ), Error ( ), Warning ( ), or Normal ( ).

You can move the pointer over the status to view more information about the event or events generated for the aggregate.

If the status of the aggregate is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event.

If the status of the aggregate is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the

administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

Aggregate

Displays the name of the aggregate.

You can move the pointer over an aggregate to view information such as the last generated event, node that contains the aggregate, RAID type, Snapshot reserve, Snapshot copies, and space allocated in the aggregate. You can also view the number of volume move operations that are currently in progress.

State

Displays the current state of the aggregate, which can be one of the following:

- Offline
  Read or write access is not allowed.
- Online
  Read and write access to volumes hosted on this aggregate is allowed.
- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
- Creating
  The aggregate is being created.
- Destroying
  The aggregate is being destroyed.
- Failed
  The aggregate cannot be brought online.
- Frozen
  The aggregate is (temporarily) not serving requests.
- Inconsistent
  The aggregate has been marked corrupted; contact technical support.
- Iron Restricted
  Diagnostic tools cannot be run on the aggregate.
- Mounting
  The aggregate is being mounted.
- Partial
  At least one disk was found for the aggregate, but two or more disks are missing.
- Quiesced
  The aggregate is quiesced.
- Quiescing
  The aggregate is being quiesced.
- Reverted
  The revert operation of the aggregate is completed.

	• Unmounted The aggregate is offline. • Unmounting The aggregate is being taken offline. • Unknown Specifies that the aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.
Node	Displays the name of the storage controller that contains the aggregate.
Flash Pool	Indicates whether or not the aggregate is a Flash Pool aggregate. The column displays either Yes or No.
Used Data Capacity	Displays the amount of space (in MB, GB, and so on) used for data in the aggregate. By default, this column is hidden.
Used Data %	Displays the percentage of space used for data in the aggregate. By default, this column is hidden.
Available Data Capacity	Displays the amount of space (in MB, GB, and so on) available for data in the aggregate.
Available Data %	Displays the percentage of space available for data in the aggregate. By default, this column is hidden.
Total Data Capacity	Displays the total data size (in MB, GB, and so on) of the aggregate.
Committed Capacity	Displays the total space (in MB, GB, and so on) committed for all the volumes in the aggregate.
RAID Type	Displays the RAID configuration type, which can be one of the following: • RAID0: All the RAID groups are of type RAID0. • RAID4: All the RAID groups are of type RAID4. • RAID-DP: All the RAID groups are of type RAID-DP. • Mixed RAID: The aggregate contains RAID groups of different RAID types (RAID0, RAID4, and RAID-DP).

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the aggregates list. You can select filters related to the Status column.

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

Related tasks

Editing aggregate threshold settings on page 101
Exporting data to CSV files on page 18

Aggregate details page

You can use the Aggregate details page to view detailed information about the selected aggregate that is monitored by Unified Manager, such as the capacity, disk information, configuration details, and events generated. You can also view information about the related objects and related alerts for that aggregate.

- *Command buttons* on page 178
- *Capacity tab* on page 178
- *Disk Information tab* on page 181
- *Configuration tab* on page 181
- *History area* on page 182
- *Events list* on page 183
- *Related Devices pane* on page 183
- *Related Alerts pane* on page 183

Command buttons

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

Actions	• Add Alert Enables you to add an alert to the selected aggregate. • Edit Thresholds Enables you to modify the threshold settings for the selected aggregate.
View Aggregates	Enables you to navigate to the Aggregates page.

Capacity tab

The Capacity tab displays detailed information about the selected aggregate, such as its capacity, thresholds, and daily growth rate. By default, the capacity events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to the node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are modified by the technical support representative, the threshold values are applied to the node root aggregate.

Capacity	Displays the data capacity graph and the Snapshot copies graph, which display capacity details about the aggregate. • Snapshot Overflow Displays the data space that is consumed by the Snapshot copies.

- Used
  Displays the space used by data in the aggregate.
- Overcommitted
  Indicates that the space in the aggregate is overcommitted.
- Warning
  Indicates that the space in the aggregate is nearly full. If this threshold is breached, the Space Nearly Full event is generated.
- Error
  Indicates that the space in the aggregate is full. If this threshold is breached, the Space Full event is generated.
- Data graph
  Displays the total data capacity and the used data capacity of the aggregate. If the aggregate is overcommitted, a flag is displayed with the overcommitted capacity.
- Snapshot Copies graph
  This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

Total Capacity Displays the total capacity in the aggregate.

Data Capacity Displays the amount of space used by the aggregate (used capacity) and the amount of available space in the aggregate (free capacity).

Snapshot Reserve Displays the used and free Snapshot capacity of the aggregate.

Overcommitted Capacity Displays the aggregate overcommitment. Aggregate overcommitment enables you to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. When thin provisioning is in use, the total size of volumes in the aggregate can exceed the total capacity of the aggregate.

> **Note:** If you have overcommitted your aggregate, you must monitor its available space carefully and add storage as required to avoid write errors due to insufficient space.

Total Cache Space Displays the total space of the solid-state disks (SSDs) that are added to a Flash Pool enabled aggregate. If you have enabled Flash Pool for an aggregate but have not added any SSDs, then the cache space is displayed as 0 KB.

> **Note:** This field is hidden if Flash Pool is disabled for an aggregate.

Aggregate Thresholds Displays the following aggregate capacity thresholds.

- Nearly Full Threshold

Specifies the percentage at which an aggregate is nearly full.

- Full Threshold
  Specifies the percentage at which an aggregate is full.
- Nearly Overcommitted Threshold
  Specifies the percentage at which an aggregate is nearly overcommitted.
- Overcommitted Threshold
  Specifies the percentage at which an aggregate is overcommitted.

Daily Growth Rate

Displays the disk space used in the aggregate if the rate of change between the last two samples continues for 24 hours.

For example, if an aggregate uses 10 GB of disk space at 2 pm and 12 GB at 6 pm, the daily growth rate (GB) for this aggregate is 2 GB.

Volume Move

Displays the number of volume move operations that are currently in progress.

- Volumes Out
  Displays the number and capacity of the volumes that are being moved out of the aggregate.
  You can click the link to view more details such as the volume name, aggregate to which the volume is moved, status of the volume move operation, and the estimated end time.
- Volumes In
  Displays the number and remaining capacity of the volumes that are being moved into the aggregate.
  You can click the link to view more details such as the volume name, aggregate from which the volume is moved, status of the volume move operation, and the estimated end time.
- Estimated used capacity after volume move
  Displays the estimated amount of used space (in percentage, and 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, the fastest daily growth rate, and the

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

The 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, the types of disks used (such as SAS, ATA, or FCAL), and the empty slots of the disks that can be added to the aggregate. 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:

RAID Details	• Type Displays the RAID type (RAID0, RAID4, RAID-DP, or Mixed RAID). • 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 disks (for example, ATA, SATA, or FCAL) in the aggregate. • Data Disks Displays the number and capacity of the data disks that are assigned to an aggregate. • Parity Disks Displays the number and capacity of the parity disks that are assigned to an aggregate.
Spare Disks	Displays the disk effective type, number, and capacity of the spare 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 the spare disks compatible with the aggregate.

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, or Mixed RAID).
- RAID Size
  Displays the size of the RAID group.
- RAID Groups
  Displays the number of RAID groups in the aggregate.
- Flash Pool
  Indicates whether or not the aggregate is a Flash Pool.

History area

The History area displays graphs that provide information about the capacity of the selected aggregate.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends—for example, if the aggregate usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Aggregate Capacity Used (%)

Displays the used capacity in the aggregate and the trend in how aggregate capacity is used based on the usage history as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Capacity Used legend, the Capacity Used graph line is hidden.

Aggregate Capacity Used vs Total Capacity

Displays the trend in how aggregate capacity is used based on the usage history, as well as the used capacity, and the total capacity as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Aggregate Capacity Used (%) vs Committed (%)	Displays the trend in how aggregate capacity is used based on the usage history, as well as the committed space as line graphs, 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 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 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 which 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 display the Critical status, then the status displayed is Critical.
Volumes	Displays the number and capacity of the volumes in the selected aggregate. The health status of the volumes is also displayed, based on the highest severity level.
Resource Pool	Displays the resource pools related to the aggregate.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected aggregate. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Related tasks

Adding an alert on page 77

Editing aggregate threshold settings on page 101

Volumes page

The Volumes page displays information about the volumes in the storage systems that are monitored and enables you to modify the volume threshold settings.

- *Command buttons* on page 184
- *Volumes Overview* on page 185
- *Volumes Protection view* on page 187
- *Filters pane* on page 190

Command buttons

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

Edit Thresholds	Displays the Edit Thresholds dialog box, which enables you to edit the threshold settings for one or more volumes.
Protect	Displays the following submenus: • SnapMirror Enables you to create a SnapMirror relationship for the selected volumes. • SnapVault Enables you to create a SnapVault relationship for the selected volumes.
Relationship	Enables you to execute the following operations: • Edit Launches the Edit Relationship dialog box which enables you to change protection policies, schedules, and maximum transfer rates for an existing protection relationship. • Abort Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer. • Quiesce Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress are completed before the relationship is quiesced. Break Breaks the relationship between the source and destination volume. After a relationship is broken, the destination is made a read-write volume. • Remove Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed.

- Resume
  Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used if one exists.
- Resynchronize
  Enables you to resynchronize a previously broken relationship.
- Initialize/Update
  Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform an update if the relationship is already initialized.
- Reverse Resync
  Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by the by the contents on the destination.

Restore

Displays the Restore dialog box, which enables you to restore directories or files from one volume at a time. If more than one volume is selected, the **Restore** button is disabled.

Export

Enables you to export the details of all the monitored FlexVol volumes to a comma-separated values (`.csv`) file. However, if you are monitoring Infinite Volumes, you cannot export the details of constituents in an Infinite Volume.

Volumes Overview

The volumes list displays, in tabular format, the properties of all the discovered volumes. You can use the column filters to customize the data that is displayed:

Status

Displays the current status of a volume. The status can be Critical (), Error (), Warning (), or Normal ().

You can move the pointer over the status to view more information about the event or events generated for the volume.

If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** link to view more information about the event.

If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

Volume	Displays the name of the volume. You can move the pointer over a volume to view information such as the qtree quota overcommitted space, status of the last volume move operation, and space allocated in the volume. You can also view the details of related objects such as the SVM to which the volume belongs, the aggregate to which the volume belongs, and all the volumes that belong to this aggregate. If an SVM with Infinite Volume is monitored, you can view details about all the three types of constituents (data, namespace, and namespace mirror) in the SVM with Infinite Volume. The constituent details include the following information: • Constituent name • State of the constituent • Name of the SVM with Infinite Volume to which the constituent belongs • Junction path of the constituent • Name of the aggregate that contains the constituent • Available, used, and total data capacity of the constituent
State	Displays the current state of the volume. The state can be Offline, Online, or Restricted.
Junction Path	Displays the path to which the volume is mounted.
Storage Virtual Machine	Displays the SVM that contains the volume.
Aggregate	Displays the name of the aggregate that contains the volume. You can view more details about the aggregate by clicking the aggregate name.
Thin Provisioned	Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.
Available Data Capacity	Displays the amount of space (in KB, MB, GB, and so on) available for data in the aggregate.
Available Data %	Displays the percentage of space available for data in the aggregate.
Used Data Capacity	Displays the amount of space (in KB, MB, GB, and so on) used for data in the volume.
Used Data %	Displays the percentage of space used for data in the volume.
Total Data Capacity	Displays the total data size (in MB, GB, and so on) of the volume.
Storage Class	Displays the storage class name. This column is displayed for Infinite Volume only.

Constituent Role	Displays the role name of the constituent. The roles can be Namespace, Data, or Namespace Mirror. This column is displayed for Infinite Volume only.
Move Status	Displays the current status of the volume move operation. The status can be In Progress, Paused, Failed, or Completed. You can move the pointer over the status to view more information about the volume move operation, such as the source, destination, operation start time, operation end time, current phase of the volume move operation which is in progress, status (in percentage), and estimated end time.

Volumes Protection view

The Protection view displays information about protection-related objects on the storage system. When you click the **Protection** button on the Volumes page, only the protection-related properties are displayed for destination volumes. When the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2, no information is displayed in some columns.

Note: If a Volumes page with Infinite Volume is monitored, protection information for constituents in the Infinite Volume are not applicable.

Relationship Status	Displays the current status of the protection relationship. The status can be either Error ( ), Warning ( ), or Normal ( ). This column is blank when a primary volume is displayed.
Transfer Status	Displays the SnapMirror status of a protection relationship. The transfer status can be one of the following:

Idle	Transfers are enabled and no transfer is in progress.
Transferring	SnapMirror transfers are enabled and a transfer is in progress.
Checking	The destination volume is undergoing a diagnostic check and no transfer is in progress.
Quiescing	A SnapMirror transfer is in progress. Additional transfers are disabled.
Quiesced	SnapMirror transfers are disabled. No transfer is in progress.
Queued	SnapMirror transfers are enabled. No transfers are in progress.
Preparing	SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.
Finalizing	SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

	Aborting SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.
Volume	Displays the name of the destination volume. You can move the pointer over a volume to view information such as the aggregate containing the volume, qtree quota overcommitted space, status of the last volume move operation, and space allocated in the volume. You can also view the details of related objects such as the SVM to which the volume belongs, the aggregate to which the volume belongs, and all the volumes that belong to this aggregate.
Storage Virtual Machine (SVM)	Displays the SVM that contains the volume.
Primary	Displays whether the volume is a primary volume or a replicated volume. If the volume is a primary volume, the column displays Yes, and protection-related columns are blank. If the volume is a replica, the column displays No.
Relationship Type	Displays the relationship type used to replicate a volume. Relationship types include SnapMirror, SnapVault, Load Sharing, Restore. This column is blank when a primary volume is displayed.
Source Volume	Displays the source volume being protected. You can view more details about the source volume by clicking the source volume name. This column is blank when a primary volume is displayed.
Lag Duration	The amount of time that the data on the mirror lags behind the source. This column is blank when a primary volume is displayed, when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2.
Lag Status	Displays the lag status for managed relationships, and for unmanaged relationships that have a schedule associated with that relationship. This column is blank when a primary volume is displayed, when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2. Lag status can be one of the following: **Error ( )** The lag duration is great than or equal to the lag error threshold. **Warning ( )** The lag duration is great than or equal to the lag warning threshold. **Normal ( )** The lag duration is within normal limits.

Last Successful Update	Display the time of the last successful data update. The time is displayed in the time zone of the Unified Manager server. This column is blank when a primary volume is displayed, when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2.
Last Transfer Duration	Displays the time taken for the last data transfer to complete. This column is blank when a primary volume is displayed, when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2.
Last Transfer Size	Displays the size, in bytes, of the last data transfer. This column is blank when a primary volume is displayed, when the monitored storage system is running a clustered Data ONTAP version earlier than 8.2, or when the protection relationships were created using a Data ONTAP version earlier than 8.2.
Relationship Health	Displays the health of the mirror. The health displays Bad if the last manual or scheduled update fails or is aborted, or if the last scheduled update is delayed. The health displays Good if the mirror is healthy and there are no issues. The health column can also remain empty if the health of a mirror is undeterminable. This column is blank when a primary volume is displayed. The column is hidden by default.
Relationship State	Displays the current state of the mirror. The state can be either Uninitialized, Snapmirrored, or Broken-off. This column is blank when a primary volume is displayed. The column is hidden by default.
Compression	Indicates whether compression is enabled on the volume. The column displays either Enabled or Disabled. This column is hidden by default.
Deduplication	Displays whether deduplication is enabled on the volume. The column displays either Enabled or Disabled. This column is hidden by default.
Volume Type	Displays which volume type is used. The volume type can be either Read-write or Data-protection. This column is hidden by default.
Cluster	Displays the cluster that contains the destination volume. You can view more details about the cluster by clicking the cluster name. This column is hidden by default.
Cluster Node	Displays the cluster node to which the volume belongs. You can view more details about the cluster node by clicking the node name. This column is hidden by default.
Source Cluster Node	Displays the source cluster node to which the source volume belongs. You can view more details about the source cluster node by clicking the source cluster node name. This column remains blank when a primary volume is displayed. This column is hidden by default.

Local Snapshot Policy	Displays the local Snapshot copy policies for the volumes listed. The default policy name is Default. This column is hidden by default.
Transfer Priority	Displays the priority of the data transfer for a volume. The priority displays either Low or Normal. The transfer priority is blank if the volume is a primary volume. This column is hidden by default.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the volumes list. You can select filters related to the Volume Status, State, and Annotation columns.

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

Related tasks

Viewing the protection inventory on page 215
Editing volume threshold settings on page 102
Exporting data to CSV files on page 18

Volume details page

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

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

- *Command buttons* on page 190
- *Capacity tab* on page 192
- *Efficiency tab* on page 195
- *Configuration tab* on page 196
- *Protection tab* on page 197
- *History area* on page 202
- *Events list* on page 203
- *Related Devices pane* on page 203
- *Related Alerts pane* on page 204
- *Annotations pane* on page 204

Command buttons

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

Actions

- Add Alert
  Enables you to add an alert to the selected volume.
- Edit Thresholds
  Enables you to modify the threshold settings for the selected volume.
- Protect
  Enables you to create either SnapMirror or SnapVault relationships for the selected volume.
- Relationship
  Enables you to execute the following protection relationship operations:
  - Edit
    Launches the Edit Relationship dialog box which enables you to change existing SnapMirror policies, schedules, and maximum transfer rates for an existing protection relationship.
  - Abort
    Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer.
  - Quiesce
    Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress will complete before the relationship is quiesced.
  - Break
    Breaks the relationship between the source and destination volume and changes the destination to a read-write volume.
  - Remove
    Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed. This operation cannot be undone.
  - Resume
    Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used if one exists.
  - Resynchronize
    Enables you to resynchronize a previously broken relationship.
  - Initialize/Update
    Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform a manual update if the relationship is already initialized.
  - Reverse Resync
    Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by

the contents on the destination, and any data that is newer than the data on the common Snapshot copy is deleted.

- Restore

  Enables you to restore data from one volume to another volume.

View Volumes Enables you to navigate to the Volumes page.

Capacity tab

The Capacity tab displays details about the selected volume, such as its capacity, threshold settings, quota capacity, and information about any volume move operation:

Capacity Displays the display capacity details 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 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 considers 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 will automatically grow in size when it is out of space.

Space Guarantee Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate. These blocks are then guaranteed to be available for writes to files in the volume. The space guarantee can be set to one of the following:

None No space guarantee is configured for the volume.

File Full size of sparsely written files, for example LUNs, is guaranteed.

Volume Full size of the volume is guaranteed.

Partial The FlexCache volume reserves space based on its size. If the FlexCache volume's size is 100 MB or more, the minimum space guarantee is set to 100 MB by default. If the FlexCache volume's size is less than 100 MB, the minimum space guarantee is set to the FlexCache volume's size. If the FlexCache volume's size is grown later, the minimum space guarantee is not incremented.

Note: The space guarantee is Partial when the volume is of type Data-Cache.

Total Capacity Displays the total capacity in the volume.

Data Capacity Displays the amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

Snapshot Reserve Displays the amount of space used by the Snapshot copies (used capacity) and amount of space available for Snapshot copies (free capacity) in the volume.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the Snapshot copies (used

capacity) and the amount of space that is available in the volume but cannot be used for making Snapshot copies (unusable capacity) because of aggregate capacity issues is displayed.

For volumes in a cluster running Data ONTAP 8.1.x, if the Snapshot used reserve is less than 1%, the **Snapshot Reserve Used** field might display a value of 0% even when there is some used data.

Volume Thresholds

Displays the following volume capacity thresholds:

- Nearly Full Threshold
  Specifies the percentage at which a volume is nearly full.
- Full Threshold
  Specifies the percentage at which a volume is full.

Other Details

- Autogrow Max Size
  Displays the maximum size up to which the FlexVol volume can automatically grow. The default value is 120% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.
- Autogrow Increment Size
  Displays the increment size using which the size of the FlexVol volume increases every time the volume is automatically grown. The default is 5% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.
- Qtree Quota Committed Capacity
  Displays the space reserved in the quotas.
- Qtree Quota Overcommitted Capacity
  Displays the amount of space that can be used before the system generates the Volume Qtree Quota Overcommitted event.
- Fractional Reserve
  Controls the size of the overwrite reserve. By default, the fractional reserve is set to 100, indicating that 100 percent of the required reserved space is reserved so that the objects are fully protected for overwrites. If the fractional reserve is less than 100 percent, the reserved space for all the space-reserved files in that volume is reduced to the fractional reserve percentage.
- Snapshot Daily Growth Rate
  Displays the change (in percentage, and 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, negative, or when there is insufficient data to calculate the growth rate.

- Snapshot Autodelete

  Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.
- Snapshot Copies

  Displays information about the Snapshot copies in the volume.

  For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

  For volumes in a cluster running Data ONTAP 8.1.x, the View link is displayed. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

Volume Move

Displays the status of either the current or the last volume move operation that was performed on the volume and other details, such as the current phase of the volume move operation which is in progress, source aggregate, destination aggregate, start time, end time, and estimated end time.

It also displays the number of volume move operations that are performed on the selected volume. You can view more information about the volume move operations by clicking the Volume Move History link.

Efficiency tab

The Efficiency tab displays information about the space saved in the volumes by using storage efficiency features such as deduplication, compression, and FlexClone volumes:

Deduplication

- Enabled

  Specifies whether deduplication is enabled or disabled on a volume.
- Space Savings

  Displays the amount of space saved (in percentage, or 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 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 KB, MB, GB, and so on) in a volume by using compression.

Configuration tab

The Configuration tab displays details about the selected volume, such as the export policy, RAID type, capacity and storage efficiency related features of the volume:

Overview

- Full Name
  Displays the full name of the volume.
- Aggregate
  Displays the name of the aggregate that contains the volume.
- Storage Virtual Machine
  Displays the name of the Storage Virtual Machine (SVM) that contains the volume.
- Junction Path
  Displays the status of the path, which can be active or inactive. The path in the SVM to which the volume is mounted is also displayed. You can click the **History** link to view the last five changes to the junction path.
- Export policy
  Displays the name of the export policy that is created for the volume. You can click the link to view details about the export policies, authentication protocols, and access enabled on the volumes that belong to the SVM.
- Type
  Displays the type of the selected volume. The volume type can be Read-write, Load-sharing, Data-Protection, Data-cache, or Temporary.
- Style
  Displays the volume style, which is FlexVol.
- RAID Type
  Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, or RAID-DP.

Capacity

- Thin Provisioning
  Displays whether thin provisioning is configured for the volume.
- Autogrow
  Displays whether the flexible volume grows in size automatically within an aggregate.
- Snapshot Autodelete
  Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.
- Quotas
  Specifies whether the quotas are enabled for the volume.

Efficiency

- Deduplication
  Specifies whether deduplication is enabled or disabled for the selected volume.
- Compression
  Specifies whether compression is enabled or disabled for the selected volume.

Protection

- Snapshot Copies
  Specifies whether automatic Snapshot copies are enabled or disabled.

Protection tab

The Protection tab displays protection details about the selected volume, such as lag information, relationship type, and topology of the relationship.

Summary Displays SnapMirror and SnapVault relationships properties for a selected volume. For any other relationship type, only the Relationship Type property is displayed. If a primary volume is selected, only the Managed and Local Snapshot copy Policy are displayed. Properties displayed for SnapMirror and SnapVault relationships include the following:

- Source Volume
  Displays the name of the selected volume's source if the selected volume is a destination.
- Lag Status
  Displays the update or transfer lag status for a protection relationship. The status can be Error, Warning, or Critical.
- Lag Duration
  Displays the time by which the data on the mirror lags behind the source.
- Last Successful Update
  Displays the date and time of the last 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.

- Relationship Capability

  Indicates the Data ONTAP capabilities available to the protection relationship. The relationship capability is either pre-8.2 or 8.2 and later. A relationship capability of pre-8.2 means that relationships have not been upgraded to Data ONTAP 8.2 on both the destination and the source clusters, and cannot take advantage of new or improved protection features available in Data ONTAP 8.2 and later. A relationship capability of 8.2 and later means the destination and source clusters are using Data ONTAP 8.2 or later, and can take advantage of improved protection features.
- Protection Service

  Displays the name of the protection service if the relationship is managed by a protection partner application.
- Relationship Type

  Displays any relationship type, including SnapMirror or SnapVault.
- Relationship State

  Displays the state of the SnapMirror or SnapVault relationship. The state can be Uninitialized, SnapMirrored, or Broken-Off. If a source volume is selected, the relationship state is not applicable and is not displayed.
- Transfer Status

  Displays the transfer status for the protection relationship. The transfer status can be one of the following:

  - Idle

    Transfers are enabled and no transfer is in progress.
  - Transferring

    SnapMirror transfers are enabled and a transfer is in progress.
  - Checking

    The destination volume is undergoing a diagnostic check and no transfer is in progress. This applies only to SnapMirror relationships that have the relationship-control-plane field set to v1.
  - Quiescing

    A SnapMirror transfer is in progress. Additional transfers are disabled.
  - Quiesced

    SnapMirror transfers are disabled. No transfer is in progress.
  - Queued

    SnapMirror transfers are enabled. No transfers are in progress.
  - Preparing

    SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.
  - Finalizing

    SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.
  - Aborting

SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

- Max Transfer Rate

  Displays the maximum transfer rate for the relationship. The maximum transfer rate can be can be a numerical value in either kilobytes per second (Kbps), Megabytes per second (Mbps), Gigabytes per second (Gbps), or Terabytes per second (TBps). If No Limit is displayed, the baseline transfer between relationships is unlimited.
- SnapMirror Policy

  Displays the protection policy for the volume. DPDefault indicates the default SnapMirror protection policy, and XDPDefault indicates the default SnapVault policy. You can click the policy name to view details associated with that policy, including the following information:

  - Transfer priority
  - Ignore access time setting
  - Tries limit
  - Comments
  - SnapMirror labels
  - Retention settings
  - Actual Snapshot copies
  - Preserve Snapshot copies
  - Retention warning threshold
  - Snapshot copies with no retention settings

  In a cascading SnapVault relationship where the source is a data protection (DP) volume, only the rule “sm_created” applies.
- Update Schedule

  Displays the SnapMirror schedule assigned to the relationship. Positioning your cursor over the information icon displays the schedule details.
- Local Snapshot Policy

  Displays the Snapshot copy policy for the volume. The policy is Default, None, or any name given to a custom policy.

Views

Displays the protection topology of the selected volume. The topology includes graphical representations of all volumes that are related to the selected volume. The selected volume is indicated by a dark gray border, and lines between volumes in the topology indicate the protection relationship type. Double lines specify a SnapMirror relationship, and a single line specifies a SnapVault relationship. The direction of the relationships in the topology are displayed from left to right, with the source of each relationship on the left and the destination on the right.

Right-clicking a volume displays a menu from which you can choose either to protect the volume or restore data to it.

Right-clicking a relationship displays a menu from which you can choose to either edit, abort, quiesce, break, remove, or resume a relationship. The menus will not display in the following instances:

- If RBAC settings do not allow this action, for example, if you have only operator privileges
- When the volume ID is unknown, for example, when you have a intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 cluster volume

Clicking another volume in the topology selects and displays information for that volume. A question mark ( ) in the upper-left corner of a volume indicates that either the volume is missing or that it has not yet been discovered. It might also indicate that the capacity information is missing. Positioning your cursor over the question mark displays additional information, including suggestions for remedial action.

The topology displays information about volume capacity, lag, Snapshot copies, and last successful data transfer if it conforms to one of several common topology templates. If a topology does not conform to one of those templates, information about volume lag and last successful data transfer is displayed in a relationship table under the topology. In that case, the highlighted row in the table indicates the selected volume, and, in the topology view, bold lines with a blue dot indicate the relationship between the selected volume and its source volume.

Topology views include the following information:

- Capacity

  Displays the total amount of capacity used by the volume. Positioning your cursor over a volume in the topology displays the current warning and critical threshold settings for that volume in the Current Threshold Settings dialog box. You can also edit the threshold settings by clicking the **Edit Thresholds** link in the Current Threshold Settings dialog box. Clearing the **Capacity** check box hides all capacity information for all volumes in the topology.
- Lag

  Displays the lag duration and the lag status of the incoming protection relationships. Clearing the **Lag** check box hides all lag information for all volumes in the topology. When the **Lag** check box is dimmed, then the lag information for the selected volume is displayed in the relationship table below the topology, as well as the lag information for all related volumes.
- Snapshot

  Displays the number of Snapshot copies available for a volume. Clearing the **Snapshot** check box hides all Snapshot copy information for all volumes in the topology. Clicking a Snapshot copy icon ( ) displays the Snapshot copy list for a volume. The Snapshot copy count displayed next to the icon is updated every 15 minutes; however, the list of Snapshot copies is updated at the time that

you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon. If you are running Data ONTAP 8.1, the Snapshot copy count is not displayed in the topology.

- Last Successful Transfer

  Displays the amount, duration, time, and date of the last successful data transfer. When the **Last Successful Transfer** check box is dimmed, then the last successful transfer information for the selected volume is displayed in the relationship table below the topology, as well as the last successful transfer information for all related volumes.

History

Displays in a graph the history of incoming SnapMirror and SnapVault protection relationships for the selected volume when you are using Data ONTAP 8.2 or later relationship capabilities. No historical data is collected for relationship capabilities earlier than Data ONTAP 8.2. There are three history graphs available: incoming relationship transfer duration, incoming relationship lag size, and incoming relationship transferred size. History information is displayed only when you select a destination volume. If you select a primary volume, the graphs are empty, and the message `No data found` is displayed.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if large amounts of data are being transferred at the same time of the day or week, or if the lag warning or lag error threshold is consistently being breached, you can take the appropriate action.

History graphs display the following information:

Relationship Transfer Duration

Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum transfer duration reached in the duration period shown in the x axis. You can view the details of specific points on the graph by positioning your cursor over the area of interest.

Relationship Lag Duration

Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum lag duration reached in the duration period shown in the x axis. The horizontal orange line on the graph depicts the lag error threshold, and the horizontal yellow line depicts the lag warning threshold. Positioning your cursor over these lines displays the threshold setting. The horizontal blue line depicts the lag duration. You can view the details for specific points on the graph by positioning your cursor over an area of interest.

Relationship Transferred Size	Displays bytes, kilobytes, megabytes, and so on, on the vertical (y) axis depending on the transfer size, and displays days, months, or years on the horizontal (x) axis depending on the selected time period. The upper value on the y axis indicates the maximum transfer size reached in the duration period shown in the x axis. You can view the details for specific points on the graph by positioning your cursor over an area of interest.

History area

The History area displays graphs that provide information about the capacity and space reservations of the selected volume.

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 Devices pane

The Related Devices pane enables you to view and navigate to the SVMs, aggregates, qtrees, LUNs, and Snapshot copies that are related to the volume:

Storage Virtual Machine	Displays the capacity and the health status of the SVM that contains the selected volume.
Aggregate	Displays the capacity and the health status of the aggregate that contains the selected volume.
Volumes in the Aggregate	Displays the number and capacity of all the volumes that belong to the parent aggregate of the selected volume. The health status of the volumes is also displayed, based on the highest severity level. For example, if an aggregate contains ten volumes, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.
Qtrees	Displays the number of qtrees that the selected volume contains and the capacity of qtrees with quota that the selected volume contains. The capacity of the qtrees with quota is displayed in relation to the volume data capacity. The health status of the qtrees is also displayed, based on the highest severity level. For example, if a volume has ten qtrees, five with Warning status and the remaining five with Critical status, then the status displayed is Critical.
NFS Exports	Displays the number and status of the NFS exports associated with the volume.
CIFS Shares	Displays the number and status of the CIFS shares.

LUNs	Displays the number and total size of all the LUNs in the selected volume. The health status of the LUNs is also displayed, based on the highest severity level.
User and Group Quotas	Displays the number and status of the user and user group quotas associated with the volume and its qtrees.
FlexClone Volumes	Displays the number and capacity of all the cloned volumes of the selected volume. The number and capacity are displayed only if the selected volume contains any cloned volumes.
Parent Volume	Displays the name and capacity of the parent volume of a selected FlexClone volume. The parent volume is displayed only if the selected volume is a FlexClone volume.

Related 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.

Annotations pane

The Annotations pane enables you to view the annotation configured for the selected volume.

Related tasks

Adding an alert on page 77
Editing volume threshold settings on page 102

Export Policy Rules dialog box

The Export Policy Rules dialog box displays details about the export policies, authentication protocols, and access enabled on the volumes that belong to the Storage Virtual Machine (SVM). You can use the filters to customize the display of information in the export policy rules list. By default, the information is sorted based on the index column.

Index	Displays the index assigned to the export policy rules. It is a unique number.
Access Protocols	Displays the protocols that are enabled for the export policy rules.
Client Match	Displays the clients that have permission to access data on the volumes that belong to the SVM.
Read Only Access	Displays the authentication protocol used to read data on the volumes that belong to the SVM.
Read Write Access	Displays the authentication protocol used to read or write data on the volumes that belong to the SVM.

Data Sources page

The Data Sources page enables you to view the statuses of data sources (clusters) that are added to the Unified Manager database. You can also view other details, such as the operation supported by the data source, state of the current operation, start and end time of the operation, and description about the operation.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

- *Command buttons* on page 205
- *Data Sources list* on page 205
- *Filters Pane* on page 206

Command buttons

The command buttons enable you to perform the following tasks for a selected data source:

Remove	Enables you to remove a data source if the discovery operation fails. This button is disabled if the data source is discovered successfully and is polling data.
Rediscover	Enables you to restart the discovery operation in case of discovery failures. In case of poll failures, you should use the **Rediscover** button from the Cluster details page.
Refresh List	Refreshes the data sources list and the properties associated with the data source.

Data sources list

The Data sources list displays, in tabular format, a list of the data sources that are added to the Unified Manager database. You can use column sorting and filtering to customize the data sources that are displayed.

Status	Displays the current discovery status of the data source. The status can be Failed ( ), Completed ( ), or In Progress ( ). You can move the pointer over the status to view more information about the event or events generated for the data source. If the status of the data source is based on a single event, you can view information such as the event name, time and date when the event was generated, the name of the administrator to whom the event is assigned, and the cause of the event. You can click **View Details** to view more information about the event.
Name	Displays the name of the data source (cluster).

Operation Displays the current operation that is supported by the data source.

The following operations are supported by the data source:

- Discovery
  Specifies the operation when the data source is being discovered.
- Poll
  Specifies the operation when the data source is successfully discovered and has started polling data.
- Deletion
  Specifies the operation when the data source (cluster) is deleted from the respective storage objects list (Clusters page).

State Displays the state of the current operation supported by the data source. The state can be Failed, Completed, or In Progress.

Start Time Displays the operation start time.

End time Displays the operation end time.

Description Displays information about the current operation and operation state of the data source. The description also provides information about when the last failed or successful operation occurred.

For example, if the operation is Poll and the operation state is Failed, the following information is displayed: `Monitoring cycle failed. Last success was 4 hours 30 minutes ago.`

Filters pane

The Filters pane enables you to set filters to customize the display of information in the data sources list.

Related tasks

Restarting cluster discovery on page 122
Removing clusters on page 123

Snapshot Copies on a Volume dialog box

You can use the Snapshot Copies on a Volume dialog box to view the list of Snapshot copies. You can delete a Snapshot copy to conserve or free disk space, or if the copy is no longer required. You

can also calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted.

List view

The list view displays, in tabular format, information about the Snapshot copies on the volume. You can use the column filters to customize the data that is displayed.

Snapshot Copy	Displays the name of the Snapshot copy.
Used Space %	Displays, in percentage, the total space used by the Snapshot copy in the volume.
Total Size	Displays the total size of the Snapshot copy.
Created Time	Displays the timestamp when the Snapshot copy was created.
Dependency	Displays the applications that are dependent on the Snapshot copy. The possible values are SnapMirror, SnapVault, SnapLock, Dump, LUNs, Vclone, and Busy.

Command buttons

The command buttons enable you to perform the following tasks:

Calculate	Enables you to calculate the space that can be reclaimed by deleting one or more Snapshot copies.
Delete Selected	Deletes one or more Snapshot copies.
Close	Closes the Snapshot Copies on a Volume dialog box.
Recalculate	Enables you to calculate the space that can be reclaimed by deleting the selected Snapshot copies. The **Recalculate** button is enabled when you make any changes in the selection of the Snapshot copies.

Related tasks

Deleting Snapshot copies on page 130
Calculating reclaimable space for Snapshot copies on page 131

Managing annotations for storage objects

You can manage storage objects in your environment by annotating them and viewing the events that correspond to them.

What annotations are

Annotations enable you to dynamically tag storage objects using user-defined rules. When you tag storage objects using 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).

OnCommand Unified Manager generates several events based on the availability, capacity, protection, and performance of storage objects in your data center. Annotations help you to filter and identify only those events that correspond to the storage objects that are tagged.

Related tasks

Creating rules to annotate storage objects on page 208
Viewing annotations on page 209
Removing annotations for storage objects on page 209

Creating rules to annotate storage objects

You can create rules to annotate clusters, volumes, and Storage Virtual Machines (SVMs) using annotations. Annotating storage objects enables you to view and manage object-related events.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Administration > Manage Annotations**.
2. In the **Auto-annotation Rules** area, click **Configure**.
3. In the **Configure Workflow Annotation Rules** dialog box, enter the rule based on your criteria.
4. Click **Validate** to validate the syntax of the rule.

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

6. Verify that the storage objects you tagged are displayed in the **Members** tab of the **Manage Annotations** page.

Related references

Manage Annotations page on page 210
Configure Workflow Annotation Rules dialog box on page 212

Viewing annotations

You can view a list of the clusters, volumes, and Storage Virtual Machines (SVMs) that are annotated so that you can manage these objects based on their priority in your environment.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Administration** > **Manage Annotations**.
2. View the annotation for your storage objects by selecting **Mission Critical**, **High**, or **Low**.

Related references

Manage Annotations page on page 210

Removing annotations for storage objects

You can remove existing annotations for storage objects when you want to change or remove the prioritization of the objects.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Administration** > **Manage Annotations**.
2. Delete the rule for configuring annotations.
3. Click **Save**.

Related references

Manage Annotations page on page 210

Understanding more about annotations

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

Description of annotation types

Annotations are categorized as three types, mission critical, high, and low. These categories are the only types that enable you to annotate storage objects based on the priority of data that they contain. Annotation types are non-editable.

Mission critical	This annotation is applied to storage objects that contain mission-critical data. For example, objects that contain production applications can be considered as mission critical.
High	This annotation is applied to storage objects that contain high priority data. For example, objects that are hosting business applications can be considered high priority.
Low	This annotation is applied to storage objects that contain low priority data. For example, objects that are on secondary storage, such as backup and mirror destinations, might be of low priority.

Description of Annotations windows and dialog boxes

You can view and manage all your annotations from the Manage Annotations page. You can also configure annotation rules for your storage objects from the Configure Annotation Rules dialog box.

Manage Annotations page

The Manage Annotations page displays a list of storage objects that you have annotated. You can also perform tasks such as viewing the annotated storage objects and configuring rules to annotate the storage objects.

- *Member Object pane* on page 211
- *Member details list* on page 211
- *Auto-annotation rules* on page 211
- *Command buttons* on page 211

Member Object pane

The Member Object pane displays the types of annotations.

Mission Critical	Displays the storage objects that are annotated as mission-critical based on the data they contain.
High	Displays the storage objects that are annotated as high based on the data they contain.
Low	Displays the storage objects that are annotated as low based on the data they contain.

Member details list

The Member details list displays, in tabular format, information about the storage objects that are annotated.

Member name	Displays the name of the storage object that is annotated. You can view the path for the location of the storage object by moving the pointer over its name.
Source Type	Displays the object type that is annotated. The type can be cluster, volume or SVM.

Auto-annotation Rules area

The Auto-annotation Rules area displays the rules that are configured to annotate storage objects.

Source Type	Displays the object type that is annotated. The type can be cluster, volume or SVM.
Condition	Displays the annotation rule that is configured for a storage object.
Apply Label to Object	Displays the annotation type that is tagged for the storage object.

Command buttons

The command button enables you to configure annotations:

Configure	Opens the Configure Workflow Annotation Rules dialog box.

Related tasks

Creating rules to annotate storage objects on page 208

Configure Workflow Annotation Rules dialog box

The Configure Workflow Annotation Rules dialog box enables you to create rules to annotate clusters, volumes, and Storage Virtual Machines (SVMs).

Rules area

You must define rules that identify the annotations tagged to a storage object.

How you define rules

You must enter the rules in the order in which you want to execute them. If the first rule's criterion is met, then objects matching that criterion are annotated based on this rule. If the criterion is not met, then the criterion for the next rule is considered, and so on. Each line lists a separate rule. You can change the order of the rules.

You can apply rules based on the name and parent attributes of a storage object. You can apply rules to annotate volumes, clusters, and SVMs based on their names or the name of their parent object. For example, a volume can have a parent SVM and an SVM can have a parent cluster, but a cluster has no parent.

Rules syntax

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

- if (*object variable.name* == *'name of the object"*) then {$*object*.data-priority = *annotation*}
  - `if` is a keyword and is in lowercase.
  - The operator is `==`.
  - `data-priority` is the category of annotations.
  - The object variable can be $VOLUME.NAME, $SVM.NAME, $CLUSTER.NAME, $VOLUME.SVM.NAME, $VOLUME.CLUSTER.NAME, or $SVM.CLUSTER.NAME.
  - Object variables and annotation conditions can be in uppercase or lowercase but must not be a combination of both. For example, $volume and $VOLUME are valid, but $Volume is not a valid variable.
  - The annotation can be mission-critical, high, or low.
- if (*object variable.name* contains *'name prefix of the object"*) then {$*object*.data-priority = *annotation*}
  - `if` is a keyword and is in lowercase.
  - The operator is `contains`.
  - `data-priority` is the category of annotations.
  - The object variable can be $VOLUME.NAME, $SVM.NAME, $CLUSTER.NAME, $VOLUME.SVM.NAME, $VOLUME.CLUSTER.NAME, or $SVM.CLUSTER.NAME.

- Object variables and annotation conditions can be in uppercase or lowercase but must not be a combination of both. For example, $volume and $VOLUME are valid, but $Volume is not a valid variable.
- The annotation can be mission-critical, high, or low.

Object variables

The following are the object variables that are used to create annotation rules:

$VOLUME	Variable attribute for the name of the volume that is annotated.
$SVM	Variable attribute for the name of the SVM that is annotated.
$CLUSTER	Variable attribute for the name of the cluster that is annotated.
$VOLUME.SVM	Variable attribute for the name of the SVM that contains all the volumes that are annotated.
$VOLUME.CLUSTER	Variable attribute for the name of the cluster that contains all the volumes that are annotated.
$SVM.CLUSTER	Variable attribute for the name of the cluster that contains all the SVMs that are annotated.

Rules syntax table based on the name and parent variables

The following table lists the syntax of the rules based on the name and parent variables.

Tip: You can copy a rule from the table based on the required criterion and paste it in the Rules area. You can then modify it to suit your requirements.

If your criterion is...	**You should use the following rule...**
To set a rule for a storage object whose name contains…	if ($VOLUME.NAME contains 'payroll') then {$VOLUME.data-priority = mission-critical} For example, all volumes whose names contain the term “payroll” are annotated as mission-critical.
To set a rule for a storage object whose name is...	if ($SVM.NAME == 'finance_svm') then {$SVM.data-priority = low} For example, all SVMs with the name “finance_svm” are annotated as low.
To set a rule for storage object whose name contains...	if ($CLUSTER.NAME contains 'cluster01') then {$CLUSTER.data-priority = high} For example, all clusters whose names contain the term “cluster01” are annotated as high.

If your criterion is...	You should use the following rule...
To set a rule for storage objects within a parent object whose name is...	if ($VOLUME.SVM.NAME == 'cloud_svm') then {$VOLUME.data-priority = mission-critical} For example, all volumes that are located in SVMs with the name “cloud_svm” are annotated as mission-critical.
To set a rule for storage objects within a parent object whose name contains...	if ($VOLUME.CLUSTER.NAME contains 'cluster01') then {$VOLUME.data-priority = high} For example, all volumes that are located in clusters whose names contain the term “cluster01” are annotated as high.
To set a rule for storage objects within a parent object whose name is...	if ($SVM.CLUSTER.NAME == 'cluster02') then {$SVM.data-priority = low} For example, all SVMs that are located in clusters with the name “cluster02” are annotated as low.

Command buttons

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

Validate Validates the syntax of the created rule. If there are errors during validation, the rule that generates the error is displayed along with an error message. The error message displays that annotation rules syntax allows only “==” and “contains”. When you click the error message, the cursor points to the corresponding error in the Rules text area.

Save Saves the rule if there are no errors and applies the annotation to the storage objects. If there are errors during validation, the rule that generates the error is displayed along with an error message.

Cancel Cancels the recent changes and closes the Configure Workflow Annotation Rules dialog box.

Related tasks

Creating rules to annotate storage objects on page 208

Managing and monitoring protection relationships

OnCommand Unified Manager enables you to create protection relationships, to monitor and troubleshoot SnapMirror and SnapVault relationships on managed clusters, and to restore data when it is overwritten or lost.

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

What resource pools are

Resource pools are groups of aggregates that are created by a storage administrator using Unified Manager to provide provisioning to partner applications for backup management.

You might pool your resources based on attributes such as performance, cost, physical location, or availability. By grouping related resources into a pool, you can treat the pool as a single unit for monitoring and provisioning. This simplifies the management of these resources and allows for a more flexible and efficient use of the storage.

Viewing the protection inventory

You can use the Volumes page to monitor your protection inventory.

Steps

1. Click **Storage** > **Volumes**.
2. Click the **Protection** button located above the volume inventory table.

   The protection inventory list is displayed.

Related tasks

Adding a user on page 316

Creating a SnapVault protection relationship from the Volumes page

You can use the Volumes page to create SnapVault relationships for one or more volumes on the same Storage Virtual Machine (SVM) to enable data backups for protection purposes on volumes running Data ONTAP 8.2 or later.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role 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 **Volumes** page, right-click a volume you want to protect and select **Protect**.

   Alternatively, to create multiple protection relationships on the same SVM, select one or more volumes in the Volumes page, and click **Protect** on the toolbar.

2. Select **SnapVault** from the menu.

   The Configure Protection dialog box is launched.

3. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary volume information.

4. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then click **Apply**.

5. Complete the **Destination Information** area and the **Relationship Settings** area in the **SnapVault** tab.

6. Click **Apply**.

   You are returned to the Volumes page.

7. Click the protection configuration job link at the top of the **Volumes** page.

   If you are creating only one protection relationship, the Job details page is displayed; however, if you are creating more than one protection relationship, a filtered list of all the jobs associated with the protection operation is displayed.

8. Do one of the following:

- If you have only one job, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
- If you have more than one job:
  a. Click a job in the jobs list.
  b. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
  c. Use the **Back** button to return to the filtered list and view another job.

Related references

Configure Protection dialog box on page 272
Advanced Secondary Settings dialog box on page 261

Creating a SnapVault protection relationship from the Volume details page

You can create a SnapVault relationship using the Volume details page so that data backups are enabled for protection purposes on volumes running Data ONTAP 8.2 or later.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Protection** tab of the **Volume details** page, right-click a volume in the topology view that you want to protect.
2. Select **Protect > SnapVault** from the menu.

   The Configure Protection dialog box is launched.
3. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary resource information.

4. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.

   You are returned to the Volume details page.
7. Click the protection configuration job link at the top of the **Volume details** page.

   The Job details page is displayed.
8. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

   When the job tasks are complete, the new relationships are displayed in the Volume details page topology view.

Related references

Configure Protection dialog box on page 272
Advanced Secondary Settings dialog box on page 261

Creating a SnapMirror protection relationship from the Volume details page

You can create a SnapMirror relationship using the Volume details page so that data replication is enabled for protection purposes. SnapMirror replication enables you to restore data from the destination volume in the event of data loss on the source for volumes running Data ONTAP 8.2 or later.

Before you begin

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

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

Steps

1. In the **Protection** tab of the **Volume details** page, right-click in the topology view the name of a volume that you want to protect.
2. Select **Protect > SnapMirror** from the menu.

   The Configure Protection dialog box is displayed.
3. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
4. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.

   You are returned to the Volume details page.
7. Click the protection configuration job link at the top of the **Volume details** page.

   The jobs tasks and details are displayed in the Job details page.
8. In the **Job details** page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
9. When the job tasks are complete, click **Back** on your browser to return to the **Volume details** page.

   The new relationship is displayed in the Volume details page topology view.

Related references

Configure Protection dialog box on page 272
Advanced Destination Settings dialog box on page 262

Creating a SnapMirror protection relationship from the Volumes page

Using the Volumes page enables you to create several SnapMirror protection relationships at one time by allowing you to select more than one volume on the same Storage Virtual Machine (SVM). The volumes must be running Data ONTAP 8.2 or later.

Before you begin

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

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Volumes** page, locate a volume you want to protect and right-click it.

   Alternatively, to create multiple protection relationships on the same SVM, select one or more volumes in the Volumes page, and click **Protect > SnapMirror** on the toolbar.

   The Configure Protection dialog box is displayed.
2. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
3. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
4. Complete the **Destination Information** area and the **Relationship Settings** area in the **SnapMirror** tab.
5. Click **Apply**.

   You are returned to the Volumes page.
6. Click the protection configuration job link at the top of the **Volumes** page.

   If you are creating only one protection relationship , the Job details page is displayed; however, if you are creating more than one protection relationship, a list of all the jobs associated with the protection operation is displayed.
7. Do one of the following:
   - If you have only one job, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
   - If you have more than one job:
     a. Click a job in the jobs list.
     b. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
     c. Use the **Back** button to return to the filtered list and view another job.

Related references

Configure Protection dialog box on page 272
Advanced Destination Settings dialog box on page 262

Configuring destination efficiency settings

You can configure destination efficiency settings such as deduplication, compression, autogrow, and space guarantee on a protection destination using the Advanced Destination Settings dialog box. You use these settings when you want to maximize space utilization on a destination or secondary volume.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

By default, efficiency settings match those of the source volume, except for compression settings in a SnapVault relationship, which are disabled by default.

Steps

1. Click either the **SnapMirror** tab or the **SnapVault** tab in the **Configure Protection** dialog box, depending on the type of relationship you are configuring.
2. Click **Advanced** in the **Destination Information** area.

   The Advanced Destination Settings dialog box is opened.
3. Enable or disable the efficiency settings for deduplication, compression, autogrow, and space guarantee, as required.
4. Click **Apply** to save your selections and return to the **Configure Protection** dialog box.

Creating SnapMirror and SnapVault schedules

You can create basic or advanced SnapMirror and SnapVault schedules to enable automatic data protection transfers on a source or primary volume so that transfers take place more frequently or less frequently, depending on how often the data changes on your volumes.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task..
- You must have already completed the Destination Information area in the Configure Protection dialog box.
- You must have set up Workflow Automation to perform this task.

Steps

1. From the **SnapMirror** tab or **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Schedule** link in the **Relationship Settings** area.

   The Create Schedule dialog box is displayed.
2. In the **Schedule Name** field, type the name you want to give to the schedule.
3. Select one of the following:
   - **Basic**
     Select if you want to create a basic interval-style schedule.
   - **Advanced**
     Select if you want to create a cron-style schedule.
4. Click **Create**.

   The new schedule is displayed in the SnapMirror Schedule or SnapVault Schedule drop-down list.

Related references

Create Schedule dialog box on page 278

Editing protection relationships from the Volumes page

You can edit existing protection relationships to change the maximum transfer rate, the protection policy, or the protection schedule. You might edit a relationship to decrease the bandwidth used for transfers, or to increase the frequency of scheduled transfers because data is changing often.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must be running Data ONTAP 8.2 or later.

About this task

The selected volumes must be protection relationship destinations. You cannot edit relationships when source volumes, load-sharing volumes, or volumes that are not the destination of a SnapMirror or SnapVault relationship are selected.

Steps

1. From the **Volumes** page, select in the volumes list one or more volumes in the same Storage Virtual Machine (SVM) for which you want to edit relationship settings, and then select **Relationship > Edit** from the toolbar.

   The Edit Relationship dialog box is displayed.

2. In the **Edit Relationship** dialog box, edit the maximum transfer rate, protection policy, or protection schedule, as needed.
3. Click **Apply**.

   The changes are applied to the selected relationships.

Editing protection relationships from the Volume details page

You can edit existing protection relationships to change the current maximum transfer rate, protection policy, or protection schedule. You might edit a relationship to decrease the bandwidth used for transfers, or to increase the frequency of scheduled transfers because data is changing often.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must be running Data ONTAP 8.2 or later.
- You must have installed and configured Workflow Automation.

About this task

The selected volumes must be protection relationship destinations. You cannot edit relationships when source volumes, load-sharing volumes, or volumes that are not the destination of a SnapMirror or SnapVault relationship are selected.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the protection relationship you want to edit and right-click it.
2. Select **Edit** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship > Edit** to edit the relationship for which you are currently viewing the details.

   The Edit Relationship dialog box is displayed.
3. In the **Edit Relationship** dialog box, edit the maximum transfer rate, protection policy, or protection schedule, as needed.
4. Click **Apply**.

   The changes are applied to the selected relationships.

Creating SnapMirror policies

You can create a SnapMirror policy to specify the SnapMirror transfer priority for protection relationships. SnapMirror policies enable you to maximize transfer efficiency from the source to the destination by assigning priorities so that lower-priority transfers are scheduled to run after normal-priority transfers.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to enable this operation.
- This task assumes that you have already completed the Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapMirror** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

   The Create SnapMirror Policy dialog box is displayed.
2. In the **Policy Name** field, type a name you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority you want to assign to the policy.
4. In the **Comment** field, enter an optional comment for the policy.
5. Click **Create**.

   The new policy is displayed in the SnapMirror Policy drop-down list.

Related references

Create SnapMirror Policy dialog box on page 279

Creating SnapVault policies

You can create a new SnapVault policy to set the priority for a SnapVault transfer. You use policies to maximize the efficiency of transfers from the primary to the secondary in a protection relationship.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to enable this operation.

- You must have already completed Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

   The SnapVault tab is displayed.
2. In the **Policy Name** field, type the name that you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority that you want to assign to the policy.
4. Optional: In the **Comment** field, enter a comment for the policy.
5. In the **Replication Label** area, add or edit a replication label, as necessary.
6. Click **Create**.

   The new policy is displayed in the Create Policy drop-down list.

Related references

Create SnapVault Policy dialog box on page 280

Aborting an active data protection transfer from the Volume details page

You can abort an active data protection transfer when you want to stop a SnapMirror replication that is in progress. You can also clear the restart checkpoint for a transfer if it is not a baseline transfer. You might abort a transfer when it conflicts with another operation, such as a volume move.

Before you begin

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

About this task

The abort action does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 or earlier cluster volume

You cannot clear the restart checkpoint for a baseline transfer.

Steps

1. In the **Protection** tab of the **Volume details** page, right-click the relationship in the topology view for the data transfer you want to abort and select **Abort**.

   The Abort Transfer dialog box is displayed.

2. If you want to clear the restart checkpoint for a transfer that is not a baseline transfer, select **Clear Checkpoints**.

3. Click **Continue**.

   The Abort Transfer dialog box is closed, and the status of the abort operation displays at the top of the Volume details page along with a link to the job details.

4. Optional: Click the **View details** link to go to the **Job details** page for additional details and to view job progress.

5. Click each job task to view its details.

6. Click the Back arrow on your browser to return to the **Volume details** page.

   The abort operation is finished when all job tasks successfully complete.

Aborting an active data protection transfer from the Volumes page

You can abort an active data protection transfer when you want to stop a SnapMirror replication that is in progress. You can also clear the restart checkpoint for transfers subsequent to the baseline transfer. You might abort a transfer when it conflicts with another operation, such as a volume move.

Before you begin

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

About this task

The abort action does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 or earlier cluster volume

You cannot clear the restart checkpoint for a baseline transfer.

Steps

1. From the **Volumes** page, right-click a volume with a relationship for which you want to abort transfers, and select **Relationship > Abort**.

   Alternatively, to abort transfers for multiple relationships, select one or more volumes, and then select **Relationship > Abort** from the toolbar.

   The Abort Transfer dialog box is displayed.
2. If you want to clear the restart checkpoint for a transfer that is not a baseline transfer, select **Clear Checkpoints**.
3. Click **Continue**.

   The Abort Transfer dialog box is closed, and the status of the abort job displays at the top of the Volumes page, along with a link to the job details.
4. Optional: Click the **View details** link to go to the **Job details** page for additional details and to view job progress.

Quiescing protection relationships from the Volume details page

You can quiesce a protection relationship to temporarily prevent data transfers from occurring. You might quiesce a relationship when you want to create a Snapshot copy of a SnapMirror destination volume that contains a database, and you want to ensure that its contents are stable during the Snapshot copy.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

About this task

The quiesce action does not display in the following instances:

- If RBAC settings do not allow this action, for example, if you have only operator privileges
- When the volume ID is unknown, for example, when you have a intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 or earlier cluster volume
- When you have not paired Workflow Automation and Unified Manager

Steps

1. In the **Protection** tab of the **Volume details** page, right-click the relationship in the topology view for the protection relationship that you want to quiesce.

2. Select **Quiesce** from the menu.
3. Click **Yes** to continue.

   The status of the quiesce job is displayed at the top of the Volume details page, along with a link to the job details.
4. Click the **View details** link to go to the **Job details** page for additional details and job progress.
5. Optional: Click the Back arrow on your browser to return to the **Volume details** page.

   The quiesce job is finished when all job tasks are successfully completed.

Breaking a SnapMirror relationship from the Volume details page

You can break a protection relationship from the Volume details page and stop data transfers between a source and destination volume in a SnapMirror relationship. You might break a relationship when you want to migrate data, for disaster recovery, or for application testing. The destination volume is changed to a read-write volume. You cannot break a SnapVault relationship.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

Steps

1. In the **Protection** tab of the **Volume details** page, select from the topology the SnapMirror relationship you want to break.
2. Right-click the destination and select **Break** from the menu.

   The Break Relationship dialog box is displayed.
3. Click **Continue** to break the relationship.
4. In the topology, verify that the relationship is broken.

Removing a protection relationship from the Volume details page

You can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

Steps

1. In the **Protection** tab of the **Volume details** page, select from the topology the SnapMirror relationship you want to remove.
2. Right-click the name of the destination and select **Remove** from the menu.

   The Remove Relationship dialog box is displayed.
3. Click **Continue** to remove the relationship.

   The relationship is removed from the Volume details page.

Resuming scheduled transfers on a quiesced relationship from the Volumes page

After you have quiesced a relationship to stop scheduled transfers from occurring, you can use **Resume** on the Volumes page to reenable scheduled transfers so that data on the source or primary volume is protected. Transfers resume from a checkpoint, if one exists, at the next scheduled transfer interval.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

About this task

You can select no more than ten quiesced relationships on which to resume transfers.

Steps

1. From the **Volumes** page, right-click a volume with a quiesced relationship, and select **Relationship > Resume**.

   Alternatively, to resume transfers for multiple relationships, select one or more volumes with quiesced relationships, and then select **Relationship > Resume** from the tool bar.

2. In the **Resume** dialog box, click **Continue**.

   You are returned to the Volume details page.

3. To view the related job tasks and to track their progress, click the job link that is displayed at the top of the **Volume details** page.

4. Do one of the following:

   - If only one job is displayed, in the Job details page click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
   - If more than one job is displayed,

     a. In the Jobs page, click the job for which you want to view the details.

     b. In the Job details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

   After the jobs complete, data transfers are resumed at the next scheduled transfer interval.

Resuming scheduled transfers on a quiesced relationship from the Volume details page

After you have quiesced a relationship to stop scheduled transfers from occurring, you can use **Resume** on the Volume details page to reenable scheduled transfers so that data on the source or primary volume is protected. Transfers resume from a checkpoint, if one exists, at the next scheduled transfer interval.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

Steps

1. In the **Protection** tab of the **Volume details** page, right-click in the topology view a quiesced relationship that you want to resume.

   Alternatively, select **Resume** from the **Actions > Relationship** menu.

2. In the **Resume** dialog box, click **Continue**.

You are returned to the Volume details page.

3. To view the related job tasks and to track their progress, click the job link that is displayed at the top of the **Volume details** page.
4. In the **Job details** page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

   After the jobs are complete, data transfers are resumed at the next scheduled transfer interval.

Initializing or updating protection relationships from the Volume details page

You can perform a first-time baseline transfer on a new protection relationship, or update a relationship if it is already initialized and you want to perform a manual, unscheduled incremental update to transfer data immediately.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must be running Data ONTAP 8.2 or later.
- You must have set up OnCommand Workflow Automation to enable this operation.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the protection relationship that you want to initialize or update, and then right-click it.
2. Select **Initialize/Update** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship > Initialize/Update** to initialize or update the relationship for which you are currently viewing the details.

   The Initialize/Update dialog box is displayed.
3. In the **Transfer Options** tab, select a transfer priority and the maximum transfer rate.
4. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.
5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
6. Click **Submit**.

   You are returned to the Initialize/Update dialog box.
7. If you selected more than one source to initialize or update, click **Default** for the next read/write source for which you want to specify an existing Snapshot copy.

You cannot select a different Snapshot copy for data protection volumes.

8. Click **Submit** to begin the initialization or update job.

   The initialization or update job is started, you are returned to the Volume details page, and a jobs link is displayed at the top of the page.

9. Optional: Click **View Jobs** on the **Volume details** page to track the status of each initialization or update job.

   A filtered list of jobs is displayed.

10. Optional: Click each job to see its details.

11. Optional: Click the Back arrow on your browser to return to the **Volume details** page.

    The initialization or update operation is finished when all job tasks successfully complete.

Initializing or updating protection relationships from the Volumes page

You can perform a first-time baseline transfer on a new protection relationship, or update a relationship if it is already initialized and you want to perform a manual, unscheduled incremental update to transfer immediately.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must be running Data ONTAP 8.2 or later.
- You must have set up OnCommand Workflow Automation to perform this operation.

Steps

1. From the **Volumes** page, right-click a volume and select **Relationship > Initialize/Update**.

   Alternatively, to initialize or update multiple relationships, select one or more volumes with relationships that you want to update or initialize, and then select **Relationship > Initialize/Update** from the tool bar.

   The Initialize/Update dialog box is displayed.

2. In the **Transfer Options** tab, select a transfer priority and the maximum transfer rate.

3. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.

4. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.

5. Click **Submit**.

You are returned to the Initialize/Update dialog box.

6. If you selected more than one source to initialize or update, click **Default** for the next source for which you want to specify an existing Snapshot copy.
7. Click **Submit** to begin the initialization or update job.

   The initialization or update job is started, you are returned to the Volumes page, and a jobs link is displayed at the top of the page.
8. Optional: Click **View Jobs** on the **Volumes** page to track the status of each initialization or update job.

   A filtered list of jobs is displayed.
9. Optional: Click each job to see its details.
10. Optional: Click the Back arrow on your browser to return to the **Volumes** page.

    The initialization or update operation is finished when all tasks successfully complete.

Resynchronizing protection relationships from the Volume details page

You can resynchronize data on a SnapMirror or SnapVault relationship that was broken and then the destination was made read/write so that data on the source matches the data on the destination. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must be running Data ONTAP 8.2 or later.
- You must have set up OnCommand Workflow Automation to perform this operation.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the protection relationship that you want to resynchronize and right-click it.
2. Select **Resynchronize** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship > Resynchronize** to resynchronize the relationship for which you are currently viewing the details.

   The Resynchronize dialog box is displayed.
3. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
4. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

The Select Source Snapshot Copy dialog box is displayed.

5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
6. Click **Submit**.

   You are returned to the Resynchronize dialog box.
7. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
8. Click **Submit** to begin the resynchronization job.

   The resynchronization job is started, you are returned to the Volume details page and a jobs link is displayed at the top of the page.
9. Optional: Click **View Jobs** on the **Volume details** page to track the status of each resynchronization job.

   A filtered list of jobs is displayed.
10. Optional: Click the Back arrow on your browser to return to the **Volume details** page.

    The resynchronization job is finished when all job tasks successfully complete.

Resynchronizing protection relationships from the Volumes page

You can resynchronize a relationship to recover from an event that disabled your source volume, or when you want to change the current source to a different volume.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must be running Data ONTAP 8.2 or later.
- You must have set up Workflow Automation to enable this operation.

Steps

1. From the **Volumes** page, right-click a volume with a quiesced relationship and select **Relationship** > **Resynchronize**.

   Alternatively, to resynchronize transfers for multiple relationships, select one or more volumes with relationships that you want to resynchronize, and then select **Relationship** > **Resynchronize** from the toolbar.
2. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
3. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

The Select Source Snapshot Copy dialog box is displayed.

4. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
5. Click **Submit**.

   You are returned to the Resynchronize dialog box.
6. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
7. Click **Submit** to begin the resynchronization job.

   The resynchronization job is started, you are returned to the Volume details page, and a jobs link is displayed at the top of the page.
8. Optional: Click **View Jobs** on the **Volume details** page to track the status of each resynchronization job.

   A filtered list of jobs is displayed.
9. Optional: Click the Back arrow on your browser to return to the **Volumes** page.

   The resynchronization operation is finished when all job tasks successfully complete.

Reversing protection relationships from the Volume details page

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- Your system must be running Data ONTAP 8.2 or later.
- Workflow Automation must be set up to perform this operation.
- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.
- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on the reverse resynchronization relationship are the same as those on the original protection relationship.
  If policies and schedules do not exist, they are created.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the SnapMirror relationship on which you want to reverse the source and destination, and right-click it.
2. Select **Reverse Resync** from the menu.

   The Reverse Resync dialog box is displayed.
3. Verify that the relationship displayed in the **Reverse Resync** dialog box is the one for which you want to perform the reverse resynchronization operation, and then click **Submit**.

   The Reverse Resync dialog box is closed and a job link is displayed at the top of the Volume details page.
4. Optional: Click **View Jobs** on the **Volume details** page to track the status of each reverse resynchronization job.

   A filtered list of jobs is displayed.
5. Optional: Click the Back arrow on your browser to return to the **Volume details** page.

   The reverse resynchronization operation is finished when all job tasks are completed successfully.

Related references

Reverse Resync dialog box on page 288

Reversing protection relationships from the Volumes 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 be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must be running Data ONTAP 8.2 or later.
- You must have set up Workflow Automation to perform this operation.

- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.
- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on reverse resynchronization relationships are the same as those on the original protection relationship.

  If policies and schedules do not exist, they are created.

Steps

1. From the **Volumes** page, click **Protection** (to the right of the **Overview** button) to display the protection inventory.
2. Right-click a volume with a SnapMirror relationship on which you want to reverse the source and destination, and then select **Relationship** > **Reverse Resync**.

   Alternatively, to reverse transfers for multiple relationships, select one or more volumes with relationships that you want to reverse, and then select **Relationship** > **Reverse Resync** from the toolbar.
3. Verify that the relationships displayed in the **Reverse Resync** dialog box are the ones for which you want to perform the reverse resync operation, and then click **Submit**.

   The reverse resync operation is started, you are returned to the Volumes page, and a jobs link is displayed at the top of the page.
4. Optional: Click **View Jobs** on the **Volumes** page to track the status of each reverse resynchronization job.

   A filtered list of jobs related to this operation is displayed.
5. Optional: Click the **Back** arrow on your browser to return to the **Volumes** page.

   The reverse resynchronization operation is finished when all job tasks successfully finish.

Related references

Reverse Resync dialog box on page 288

Restoring data using the Volumes page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Volumes page. Restoring data from volumes that have a version of Data ONTAP that is earlier than 8.2 is not supported.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

You cannot restore NTFS file streams.

The restore option is not available in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Volumes** page, select a volume from which you want to restore data.
2. From the toolbar, click **Restore**.

   The Restore dialog box is displayed.
3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
4. Select the items you want to restore.

   You can restore the entire volume, or you can specify folders and files you want to restore.
5. Select the location to which you want the selected items restored; either **Original Location** or **Alternate Location**.
6. Click **Restore**.

   The restore process begins.

Related references

Restore dialog box on page 263

Restoring data using the Volume details page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Volume details page. Restoring data from volumes that have a version of Data ONTAP that is earlier than 8.2 is not supported.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

You cannot restore NTFS file streams.

The restore option is not available in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Protection** tab of the **Volume details** page, right-click in the topology view the name of the volume that you want to restore.
2. Select **Restore** from the menu.

   Alternatively, select **Restore** from the **Actions** menu to protect the current volume for which you are viewing the details.

   The Restore dialog box is displayed.
3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
4. Select the items you want to restore.

   You can restore the entire volume, or you can specify folders and files you want to restore.
5. Select the location to which you want the selected items restored: either **Original Location** or **Alternate Existing Location**.
6. If you select an alternate existing location, do one of the following:
   - In the Restore Path text field, type the path of the location to which you want to restore the data and then click **Select Directory**.
   - Click **Browse** to launch the Browse Directories dialog box and complete the following steps:

     a. Select the cluster, SVM, and volume to which you want to restore.

   b. In the Name table, select a directory name.
   c. Click **Select Directory**.

7. Click **Restore**.

   The restore process begins.

Related references

Restore dialog box on page 263

Creating resource pools

You can use the Create Resource Pool dialog box to group aggregates for provisioning purposes.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

Resource pools can contain aggregates from different clusters, but the same aggregate cannot belong to different resource pools.

Steps

1. Click **Storage** > **Resource Pools**.
2. On the **Resource Pools** page, click **Create**.

   The Create Resource Pool dialog box is displayed.
3. Follow the instructions in the dialog box to provide a name and description and to add aggregates as members to the resource pool you want to create.

Related tasks

Adding a user on page 316

Related references

Create Resource Pool dialog box on page 253

Editing resource pools

You can edit an existing resource pool when you want to change the resource pool name and the description.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

The **Edit** button is enabled only when one resource pool is selected. If more than one resource pool is selected, the **Edit** button is disabled.

Steps

1. Click **Storage** > **Resource Pools**.
2. Select one resource pool from the list.
3. Click **Edit**.

   The Edit Resource Pool window is displayed.
4. Edit the resource pool name and description as needed.
5. Click **Save**.

   The new name and description are displayed in the resource pool list.

Related tasks

Adding a user on page 316

Viewing resource pools inventory

You can use the Resource Pools page to view the resource pool inventory and to monitor the remaining capacity for each resource pool.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Step

1. Click **Storage** > **Resource Pools**.

The resource pool inventory is displayed.

Related tasks

Adding a user on page 316

Related references

Resource Pools page on page 251

Adding resource pool members

A resource pool consists of a number of member aggregates. You can add aggregates to existing resource pools to increase the amount of space available for secondary volume provisioning.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

You can add no more than 200 aggregates to a resource pool at one time. Aggregates shown in the Aggregates dialog box do not belong to any other resource pool.

Steps

1. Click **Storage > Resource Pools**.
2. Select a resource pool from the **Resource Pools** list.

   The resource pool members are displayed in the area below the resource pool list.
3. In the resource pool member area, click **Add**.

   The Aggregates dialog box is displayed.
4. Select one or more aggregates.
5. Click **Add**.

   The dialog box is closed and the aggregates are displayed in the member list for the selected resource pool.

Related tasks

Adding a user on page 316

Removing aggregates from resource pools

You can remove aggregates from an existing resource pool: for example, when you want to use an aggregate for some other purpose.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

Resource pool members are displayed only when a resource pool is selected.

Steps

1. Click **Storage** > **Resource Pools**.
2. Select the resource pool from which you want to remove member aggregates.

   The list of member aggregates is displayed in the Members pane.
3. Select one or more aggregates.

   The **Remove** button is enabled.
4. Click **Remove.**

   A warning dialog box is displayed.
5. Click **Yes** to continue.

   The selected aggregates are removed from the Members pane.

Related tasks

Adding a user on page 316

Deleting resource pools

You can delete resource pools when they are no longer needed. For example, you might want to redistribute the member aggregates from one resource pool to several other resource pools, making the original resource pool obsolete.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

The **Delete** button is enabled only when at least one resource pool is selected.

Steps

1. Click **Storage** > **Resource Pools**.
2. Select the resource pool you want to delete.
3. Click **Delete**.

   The resource pool is removed from the resource pool list and its aggregates are removed from the members list.

Related tasks

Adding a user on page 316

Related references

Resource Pools page on page 251

Understanding SVM associations

Storage Virtual Machine (SVM) associations are mappings from a source SVM to a destination SVM that are used by partner applications for resource selection and secondary volume provisioning.

Associations are always created between the primary SVM and the destination regardless of whether the destination is a secondary or a tertiary destination. You cannot associate a secondary destination SVM with a tertiary destination SVM.

You can associate SVMs in two ways:

- Associate any SVM

  You can create an association between any primary SVM source to one or more destination SVMs. This means that all existing SVM that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVMs. For example, you might want applications from several different sources at different locations backed up to one or more destination SVMs in one location.
- Associate a particular SVM

  You can create an association of a specific source SVM with one or more destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, you can choose this option to associate a specific SVM source to a specific SVM destination that is assigned to only that client.

Storage Virtual Machine (SVM) and resource pool requirements to support storage services

You can better ensure conformance in partner applications if you observe some SVM association and resource pool requirements that are specific to storage services: for example, when you associate Storage Virtual Machines (SVMs) and create resource pools in Unified Manager to support a protection topology in a storage service provided by a partner application.

Some applications partner with the Unified Manager server to provide services that automatically configure and execute SnapMirror or SnapVault backup protection between source volumes and protection volumes in secondary or tertiary locations. To support these protection storage services, you must use Unified Manager to configure the necessary SVM associations and resource pools.

To support storage service single-hop or cascaded protection, including replication from a SnapMirror source or SnapVault primary volume to either destination SnapMirror or to SnapVault backup volumes that reside in secondary or tertiary locations, observe the following requirements:

- SVM associations must be configured between the SVM containing the SnapMirror source or SnapVault primary volume and any SVM on which either a secondary volume or a tertiary volume resides.
  - For example, to support a protection topology in which source volume Vol_A resides on SVM_1, and SnapMirror secondary destination volume Vol_B resides on SVM_2, and tertiary SnapVault backup volume Vol_ C resides on SVM_3, you must use the Unified Manager web UI to configure a SnapMirror association between SVM_1 and SVM_2 and a SnapVault backup association between SVM_1 and SVM_3.
    In this example, any SnapMirror association or SnapVault backup association between SVM_2 and SVM_3 is not necessary and is not used.
  - To support a protection topology in which both source volume Vol_A and SnapMirror destination volume Vol_B reside on SVM_1, you must configure a SnapMirror association between SVM_1 and Vserver_1.
- The resource pools must include cluster aggregate resources available to the associated SVMs.
  You configure resource pools through the Unified Manager web UI and then assign through the partner application the storage service secondary target and tertiary target nodes.

Related tasks

Creating resource pools on page 240
Creating Storage Virtual Machine (SVM) associations on page 245

Creating Storage Virtual Machine (SVM) associations

The Create Storage Virtual Machine Associations wizard enables partner protection applications to associate source Storage Virtual Machine (SVM) with destination SVM for use with SnapMirror and

SnapVault relationships. Partner applications use these associations at the time of initial provisioning of destination volumes to determine which resources to select.

Before you begin

- The SVM you are associating must already exist.
- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

For any source SVM and relationship type, you can choose only one destination SVM on each destination cluster.

Changing associations using the delete and create functions affects only future provisioning operations. It does not move existing destination volumes.

Steps

1. Click **Storage > Storage Virtual Machine Associations**.
2. Click **Create**.

   The Create Storage Virtual Machine Associations wizard is launched.
3. Select one of the following:
   - **Associate Any Storage Virtual Machine**
     Choose this option when you want to create an association between any primary SVM source to one or more destination SVM. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVM. For example, you might want applications from several different sources at different locations backed up to one or more destination SVM in one location.
   - **Select Particular Storage Virtual Machine**
     Choose this option when you want to select a specific source SVM associated with one or more destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, choose this option to associate a specific SVM source to a specific SVM destination that is assigned only to that client.
4. Select one or both of the protection relationship types you want to create:
   - **SnapMirror**
   - **SnapVault**
5. Click **Next**.
6. Select one or more SVM protection destination.
7. Click **Finish**.

Related tasks

Adding a user on page 316

Viewing SVM associations

You can use the Storage Virtual Machine Associations page to view existing Storage Virtual Machine (SVM) associations and their properties and to determine if additional SVM associations are required.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Step

1. Click **Storage > Storage Virtual Machine Associations**.

   The list of SVM associations and their properties is displayed.

Related tasks

Adding a user on page 316

Deleting SVM associations

You can delete Storage Virtual Machine (SVM) associations for partner applications to remove the secondary provisioning relationship between source and destination SVMs; for example, you might do this when the destination SVM is full and you want to create a new SVM protection association.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

The **Delete** button is disabled until at least one SVM association is selected. Changing associations using the delete and create functions affects only future provisioning operations; it does not move existing destination volumes.

Steps

1. Click **Storage > Storage Virtual Machine Associations**.
2. Select at least one SVM association.

The **Delete** button is enabled.

3. Click **Delete.**

   A warning dialog box is displayed.

4. Click **Yes** to continue.

   The selected SVM association is removed from the list.

Related tasks

Adding a user on page 316

What jobs are

A job is a series of tasks that you can monitor using Unified Manager. Viewing jobs and their associated tasks enables you to determine if a they have completed successfully.

Jobs are initiated when you create SnapMirror and SnapVault relationships, when you perform any relationship operation (break, edit, quiesce, remove, resume, resynchronize, and reverse resync), when you perform data restoration tasks, when you log in to a cluster, and so on.

When you initiate a job, you can use the Jobs page and the Job details page to monitor the job and the progress of the associated job tasks.

Monitoring jobs

You can use the Jobs page to monitor job status and to view job properties such as storage service type, state, submitted time, and completed time to determine whether or not a job has successfully completed.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Jobs**.

   The Jobs page is displayed.

2. View the **Status** column to determine the status of those jobs currently running.

3. Click on a job name to view details about that particular job.

   The Job details page is displayed.

Related tasks

Adding a user on page 316

Viewing job details

After you start a job, you can track its progress from the Job details page and monitor the associated tasks for possible errors.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Jobs**.
2. From the **Jobs** page, click a job name in the **Name** column to display the list of tasks associated with the job.
3. Click on a task to display additional information in the **Task Details** pane and the **Task Messages** pane to the right of the task list.

Related tasks

Adding a user on page 316

Aborting jobs

You can use the Jobs page to abort a job if it is taking too long to finish, is encountering too many errors, or is no longer needed. You can abort a job only if its status and type allow it. You can abort any running job.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Jobs**.
2. From the list of jobs, select one job.
3. Click **Abort**.
4. At the confirmation prompt, click **Yes** to abort the selected job.

Related tasks

Adding a user on page 316

Retrying a failed protection job

After you have taken measures to fix a failed protection job, you can use **Retry** to run the job again. Retrying a job creates a new job using the original job ID.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

You can retry only one failed job at a time. Selecting more than one job disables the **Retry** button. Only jobs of the type Protection Configuration and Protection Relationship Operation can be retried.

Steps

1. Click **Jobs**.
2. From the list of jobs, select a single failed Protection Configuration or Protection Relationship Operation type job.

   The **Retry** button is enabled.
3. Click **Retry**.

   The job is restarted.

Related references

Job details page on page 266

Description of Protection relationships window and dialog boxes

You can view and manage protection-related details such as resource pools, Storage Virtual Machine (SVM) associations, and protection jobs. You can use the appropriate Setup Options dialog box to configure global threshold values for aggregates, volumes, and relationships.

Resource Pools page

The Resource Pools page displays existing resource pools and their members, and enables you to create, monitor, and manage resource pools for provisioning purposes.

- *Command buttons* on page 251
- *Resource Pools list* on page 251
- *Members list command buttons* on page 251
- *Members list* on page 251

Command buttons

The command buttons enable you to perform the following tasks:

Create	Launches the Create Resource Pool dialog box.
Edit	Enables you to edit the name and description of the resource pools that you create.
Delete	Enables you to delete one or more resource pools.

Resource Pools list

The resource pools list displays, in tabular format, the properties of existing resource pools.

Resource Pool	Displays the name of the resource pool.
Description	Describes the resource pool.
Total Capacity	Displays the total capacity (in MB, GB, and so on) of the resource pool.
Used Capacity	Displays the amount of space (in MB, GB, and so on) used in the resource pool.
Available Capacity	Displays the amount of space (in MB, GB, and so on) available in the resource pool.
Used %	Displays the percentage of space used in the resource pool.

Members list command buttons

The members list command buttons enable you to perform the following tasks:

Add	Enables you to add members to the resource pool.
Delete	Enables you to delete one or more members from the resource pool.

Members list

The members list displays, in tabular format, the resource pool members and their properties when a resource pool is selected.

Status

Displays the current status of the aggregate. Status can be Critical (), Error (), Warning (), or Normal ().

Aggregate Name

Displays the name of the member aggregate.

State

Displays the current state of the aggregate, which can be one of the following:

- Offline
  Read or write access is not allowed.
- Online
  Read and write access to volumes hosted on this aggregate is allowed.
- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
- Creating
  The aggregate is being created.
- Destroying
  The aggregate is being destroyed.
- Failed
  The aggregate cannot be brought online.
- Frozen
  The aggregate is (temporarily) not serving requests.
- Inconsistent
  The aggregate has been marked corrupted; contact technical support.
- Iron Restricted
  Diagnostic tools cannot be run on the aggregate.
- Mounting
  The aggregate is in the process of mounting.
- Partial
  At least one disk was found for the aggregate, but two or more disks are missing.
- Quiescing
  The aggregate is being quiesced.
- Quiesced
  The aggregate is quiesced.
- Reverted
  The revert of an aggregate is completed.
- Unmounted
  The aggregate has been unmounted.
- Unmounting
  The aggregate is being taken offline.

- Unknown
  Specifies that the aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

By default, this column is hidden.

Cluster	Displays the name of the cluster to which the aggregate belongs.
Node	Displays the name of the node on which the aggregate resides.
Total Capacity	Displays the total capacity (in MB, GB, and so on) of the aggregate.
Used Capacity	Displays the amount of space (in MB, GB, and so on) used in the aggregate.
Available Capacity	Displays the amount of space (in MB, GB, and so on) available in the aggregate.
Used %	Displays the percentage of space used in the aggregate.
Disk Type	Displays the RAID configuration type, which can be one of the following: • RAID0: All the RAID groups are of type RAID0. • RAID4: All the RAID groups are of type RAID4. • RAID-DP: All the RAID groups are of type RAID-DP. • Mixed RAID: The aggregate contains RAID groups of different RAID types (RAID0, RAID4, and RAID-DP). By default, this column is hidden.

Related tasks

Viewing resource pools inventory on page 241
Deleting resource pools on page 243
Creating resource pools on page 240
Editing resource pools on page 241
Adding resource pool members on page 242
Removing aggregates from resource pools on page 243

Create Resource Pool dialog box

You can use the Create Resource Pool dialog box to name and describe a new resource pool and to add aggregates to and delete aggregates from that resource pool.

The text boxes enable you to add the following information to create a resource pool:

Resource Pool Name

Enables you to specify a resource pool name.

Description

Enables you to describe a resource pool.

Members

Displays the members of the resource pool. You can also add and delete members.

Command buttons

The command buttons enable you to perform the following tasks:

Add	Opens the Aggregates dialog box so that you can add aggregates from a specific cluster to the resource pool. You can add aggregates from different clusters, but the same aggregates cannot be added to more than one resource pool. Aggregates belonging to clusters that are running Data ONTAP 8.1 are not displayed in the Aggregates dialog box.
Remove	Enables you to remove selected aggregates from the resource pool.
Create	Creates the resource pool. This button is not enabled until information has been entered in the Resource Pool Name or Description fields.
Cancel	Discards the changes and closes the Create Resource Pool dialog box.

Related concepts

Storage Virtual Machine (SVM) and resource pool requirements to support storage services on page 245

Related tasks

Creating resource pools on page 240
Adding resource pool members on page 242

Edit Resource Pool dialog box

You can use the Edit Resource Pool dialog box to change the name and description of an existing resource pool. For example, if the original name and description is inaccurate or incorrect, you can change them so they are more precise.

Text boxes

The text boxes enable you to change the following information for the selected resource pool:

Resource Pool Name	Enables you to enter a new name.
Description	Enables you to enter a new description.

Command buttons

The command buttons enable you to perform the following tasks:

Save	Saves the changes to the resource pool name and description.
Cancel	Discards the changes and closes the Edit Resource Pool dialog box.

Related tasks

Editing resource pools on page 241

Aggregates dialog box

You can use the Aggregates dialog box to select aggregates that you want to add to your resource pool.

- *Command buttons* on page 255
- *Aggregates list* on page 255

Command buttons

The command buttons enable you to perform the following tasks:

Add	Adds the selected aggregates to the resource pool. The Add button is not enabled until at least one aggregate is selected.
Cancel	Discards the changes and closes the Aggregates dialog box.

Aggregates list

The Aggregates list displays in a table the names and properties of monitored aggregates. If the monitored aggregates belong to a storage system running Data ONTAP 8.1, the aggregates on that storage system are not displayed in the aggregates list.

Status	Displays the current status of a volume. The status can be Critical ( ), Error ( ), Warning ( ), or Normal ( ). You can move the pointer over the status to view more information about the event or events generated for the volume.
Aggregate Name	Displays the name of the aggregate.
State	Displays the current state of the aggregate. The state can be Offline, Restricted, Online, or Unknown. • Offline Read or write access is not allowed.

- Restricted
  Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
- Online
  Read and write access to volumes hosted on this aggregate is allowed.
- Creating
  The aggregate is being created.
- Destroying
  The aggregate is being destroyed.
- Failed
  The aggregate cannot be brought online.
- Frozen
  The aggregate is (temporarily) not serving requests.
- Inconsistent
  The aggregate has been marked corrupted; contact technical support.
- Inron Restricted
  Diagnostic tools cannot be run on the aggregate.
- Mounting
  The aggregate is being mounted.
- Partial
  At least one disk was found for the aggregate, but two or more disks are missing.
- Quiescing
  The aggregate is being quiesced.
- Quiesced
  The aggregate is quiesced.
- Reverted
  The revert of an aggregate is completed.
- Unmounted
  The aggregate is offline.
- Unmounting
  The aggregate is being taken offline.
- Unknown
  Specifies that the aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

Cluster Displays the name of the cluster on which the aggregate resides.

Node Displays the name of the storage controller that contains the aggregate.

Total Capacity Displays the total data size (in MB, GB, and so on) of the aggregate. By default, this column is hidden.

Committed Capacity	Displays the total space (in MB, GB, and so on) committed for all the volumes in the aggregate. By default, this column is hidden.
Used Capacity	Displays the amount of space (in MB, GB, and so on) used in the aggregate.
Available Capacity	Displays the amount of space (in MB, GB, and so on) available for data in the aggregate. By default, this column is hidden.
Available %	Displays the percentage of space available for data in the aggregate. By default, this column is hidden.
Used %	Displays the percentage of space used by data in the aggregate.
RAID Type	Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, or Mixed RAID.

Related tasks

Adding resource pool members on page 242

Storage Virtual Machine Associations page

The Storage Virtual Machine Associations page enables you to view existing Storage Virtual Machine (SVM) associations between source and destination SVMs and to create new SVM associations for use by partner applications to create SnapMirror and SnapVault relationships.

Command buttons

The command buttons enable you to perform the following tasks:

Create	Opens the Create Storage Virtual Machine Associations wizard.
Delete	Enables you to delete the selected SVM associations.

Storage Virtual Machine (SVM) Associations list

The Storage Virtual Machine Associations list displays in a table the source and destination SVM associations that have been created and the type of protection relationship allowed for each association.

Source Storage Virtual Machine	Displays the name of the source SVM.
Source Cluster	Displays the name of the source cluster.
Destination Storage Virtual Machine	Displays the name of the destination SVM.
Destination Cluster	Displays the name of the destination cluster.

Type	Displays the type of protection relationship. Relationship types are either SnapMirror or SnapVault.

Related tasks

Viewing SVM associations on page 247
Creating Storage Virtual Machine (SVM) associations on page 245
Deleting SVM associations on page 247

Create Storage Virtual Machine Associations wizard

The Create Storage Virtual Machine Associations wizard enables you to associate source and destination Storage Virtual Machines (SVMs) for use in SnapMirror and SnapVault protection relationships.

Select Source SVM

The Select Source Storage Virtual Machine panel enables you to select the source or primary SVM in the SVM association.

Associate Any SVM	Enables you to create an association between any SVM source or primary to one or more destination or secondary SVM. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination or secondary SVM. For example, you might want applications from several different sources at different locations backed up to one or more destination SVM in one location.
Select Particular SVM	Enables you to associate a specific source SVM with one or more destination SVM. For example, if you are providing storage services to many clients whose data must be separate from one another, choose this option to associate a specific SVM source to a specific SVM destination that is assigned only to that client. • Storage Virtual Machine Lists the names of the available source SVMs • Cluster Lists the clusters on which each SVM resides
Allow these kinds of relationships	Enables you to select the relationship type for the association: • SnapMirror Specifies a SnapMirror relationship as the association type. Selecting this option enables data replication from the selected sources to the selected destinations. • SnapVault

Specifies a SnapVault relationship as the association type. Selecting this option enables backups from the selected primary locations to the selected secondary locations.

Select Protection Destinations

The Select Protection Destinations panel of the Create Storage Virtual Machine Associations wizard enables you to select where to copy or replicate the data. You can create an association on only one destination SVM per cluster.

Command buttons

The command buttons enable you to perform the following tasks:

Next	Advances you to the next page in the wizard.
Back	Returns you to the previous page in the wizard.
Finish	Applies your selections and creates the association.
Cancel	Discards the selections and closes the Create Storage Virtual Machine Associations wizard.

Jobs page

The Jobs page enables you to view the current status and other information about all partner application protection jobs that are currently running, as well as jobs that have completed. You can use this information to see which jobs are still running and whether a job has succeeded or failed.

- *Command buttons* on page 259
- *Jobs list* on page 260

Command buttons

The command buttons enable you to perform the following tasks:

Abort	Aborts the selected job. This option is available only if the selected job is running.
Retry	Restarts a failed job of type Protection Configuration or Protection Relationship Operation. You can retry only one failed job at a time. If more than one failed job is selected, the **Retry** button is disabled. You cannot retry failed storage service jobs.
Refresh	Refreshes the list of jobs and the information associated with them.

Jobs list

The Jobs list displays, in tabular format, a list of the jobs that are in progress. By default, the list displays only the jobs generated within the past week. You can use column sorting and filtering to customize which jobs are displayed.

Status	Displays the current status of a job. The status can be Error ( ) or Normal ( ).
Job Id	Displays the identification number of the job. By default, this column is hidden. The job identification number is unique and is assigned by the server when it starts the job. You can search for a particular job by entering the job identification number in the text box provided by the column filter.
Name	Displays the name of the job.
Type	Displays the storage service job type. The job types are as follows: **Cleanup** — The job is cleaning up storage service member artifacts that are no longer needed for restore purposes. **Conform** — The job is checking the configuration of storage service members to ensure that they conform. **Destroy** — The job is destroying a storage service. **Import** — The job is importing unmanaged storage objects into an existing storage service. **Modify** — The job is modifying attributes of an existing storage service. **Restore** — A restore job is running. **Subscribe** — The job is subscribing members to a storage service. **Unsubscribe** — The job is unsubscribing members from a storage service. **Update** — A protection update job is running.
State	Displays the running state of the job. State options are as follows: **Aborted** — The job has been aborted. **Aborting** — The job is in the process of aborting. **Completed** — The job has finished. **Running** — The job is running.
Submitted Time	The time the job was submitted.
Duration	The amount of time the job took to complete. This column is displayed by default.

Completed Time Displays the time the job finished. By default, this column is hidden.

Related tasks

Monitoring jobs on page 248
Viewing job details on page 249
Aborting jobs on page 249

Advanced Secondary Settings dialog box

You can enable space-related settings on a secondary volume to maximize the amount of data being stored, including the following: deduplication, data compression, autogrow, and space guarantee. You might use the Advanced Secondary Settings dialog box when want to change enable or disable the current settings.

Enable Deduplication Enables deduplication on the secondary volume in a SnapVault relationship so that duplicate data blocks are eliminated to achieve space savings. You might use deduplication is when space savings are at least 10 percent and when data overwrite rate is not rapid. Deduplication is often used for virtualized environments, file shares, and backup data. This setting is disabled by default. When enabled, this operation is initiated after each transfer.

- Enable Compression
  Enables transparent data compression. You might use compression when space savings are at least 10 percent, when the potential overhead is acceptable, and when there are sufficient system resources for compression to complete during non-peak hours. In a SnapVault relationship, this setting is disabled by default. Compression is available only when deduplication is selected.
  - Compress Inline
    Enables immediate space savings by compressing data before writing data to disk. You might use inline compression when your system has no more than 50 percent utilization during peak hours, and when the system can accommodate new writes and additional CPU during peak hours. This setting is available only when “Enable Compression” is selected.

For more information on deduplication and compression, see TR-3966 at *http://www.netapp.com/us/media/tr-3966.pdf*.

Enable Autogrow Enables you to automatically grow the destination volume when the free space percentage is below specified the threshold, as long as space is available on the associated aggregate.

Maximum Size Sets the maximum percentage to which a volume can grow. The default is 20 percent greater than the source volume size. A volume does not grow

automatically if the current size is greater than or equal to the maximum autogrow percentage. This field is enabled only when the autogrow setting is enabled.

Increment Size The percentage increment by which the volume automatically grows before reaching the maximum percentage of the source volume.

Space Guarantee Ensures that enough space is allocated on the secondary volume so that data transfers always succeed. The space guarantee setting can be one of the following:

- File
- Volume
- None

For example, you might have a 200 GB volume that contains files totaling 50 GB; however, those files hold only 10 GB of data. Volume guarantee allocates 200 GB to the destination volume, regardless of contents on the source. File guarantee allocates 50 GB to ensure that enough space is reserved for files on the source; selecting None is this scenario means that only 10 GB is allocated on the destination for the actual space used by file data on the source.

The space guarantee is set to Volume by default.

Command buttons

The command buttons enable you to perform the following tasks:

Apply Saves the selected efficiency settings and applies them when you click **Apply** in the Configure Protection dialog box.

Cancel Discards your selections and closes the Advanced Destination Settings dialog box.

Related references

Configure Protection dialog box on page 272

Advanced Destination Settings dialog box

You can use the Advanced Destination Settings dialog box to enable space guarantee settings on a destination volume. You might select advanced settings when space guarantee is disabled on the source, but you want it enabled on the destination. The settings for deduplication, compression, and autogrow in a SnapMirror relationship are inherited from the source volume and cannot be changed.

Space Guarantee

Ensures that enough space is allocated on the destination volume so that data transfers always succeed. The space guarantee setting can be one of the following:

- File
- Volume
- None

For example, you might have a 200-GB volume that contains files totaling 50 GB; however, those files hold only 10 GB of data. Volume guarantee allocates 200 GB to the destination volume, regardless of contents on the source. File guarantee allocates 50 GB to ensure that enough space is reserved for source files on the destination; selecting **None** is this scenario means that only 10 GB is allocated on the destination for the actual space used by file data on the source.

The space guarantee is set to Volume by default.

Related references

Configure Protection dialog box on page 272

Restore dialog box

You can use the Restore dialog box to restore data to a volume from a specific Snapshot copy. Restoring data on volumes using a Data ONTAP version earlier than 8.2 is not supported.

Restore from

The Restore from area enables you to specify from where you want to restore data.

Volume	Specifies the volume from which you want to restore data. By default, the volume on which you initiated the restore action is selected, or you can select a different volume from the drop-down list that contains all the volumes with protection relationships to the volume on which you initiated the restore action.
Snapshot copy	Specifies which Snapshot copy you want to use to restore data. By default, the most recent Snapshot copy is selected. You can also select a different Snapshot copy from the drop-down list. The Snapshot copy list changes according to which volume is selected.

Select items to restore

The Select items to restore area enables you to select either the entire volume or specific files and folders you want to restore. You can select a maximum of 10 files, folders, or a combination of both. When the maximum number of items is selected, the item selection check boxes are disabled.

Path field	Displays the path to the data you want to restore. You can either navigate to the folder and files you want to restore, or you can type the path. This field is empty until you select or type a path. Clicking after you have chosen a path moves you up one level in the directory structure.
Folders and files list	Displays the contents of the path you entered. By default, the root folder is initially displayed. Clicking a folder name displays the contents of the folder.

You can select items to restore as follows:

- When you enter the path with a particular file name specified in the path field, the specified file is displayed in the Folders and Files.
- When you enter a path without specifying a particular file, the contents of the folder are displayed in the Folders and Files list, and you can select up to 10 files, folders, or a combination of both to restore. However, if a folder contains more than 995 items, a message displays to indicate there are too many items to display and if you proceed with the operation, all items in the specified folder are restored.

Note: You cannot restore NTFS file streams.

Restore to

The Restore to area enables you to specify where you want to restore the data.

Original Location in *`Volume_Name`*	Restores the selected data to the directory on the source from which the data was originally backed up.
Alternate Location	Restores the selected data to a new location: • Restore Path Specifies an alternate path for restoring the selected data. The path must already exist. You can use the **Browse** button to navigate to the location where you want the data restored, or you can enter the path manually using the format cluster://svm/volume/path. • Preserve directory hierarchy When checked, preserves the structure of the original file or directory. For example, if the source is /A/B/C/myFile.txt and the destination is /X/Y/Z, Unified Manager restores the data using the following directory structure on the destination: /X/Y/Z/A/B/C/myFile.txt.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel	Discards your selections and closes the Restore dialog box.
Restore	Applies your selections and begins the restore process.

Related tasks

Restoring data using the Volumes page on page 238
Restoring data using the Volume details page on page 239

Related references

Browse Directories dialog box on page 265

Browse Directories dialog box

You can use the Browse Directories dialog box when you want to restore data to a directory on a cluster and Storage Virtual Machine (SVM) that is different from the original source. The original source cluster and volume are selected by default.

The Browse Directories dialog box enables you to select the cluster, SVM, volume, and directory path to which you want data restored.

Cluster	Lists the available cluster destinations to which you can restore. By default, the cluster of the original source volume is selected.
SVM drop-down list	Lists the available SVMs available for the selected cluster. By default, the SVM of the original source volume is selected.
Volume	Lists all of the read/write volumes in a selected SVM. You can filter the volumes by name and by space available. The volume with the most space is listed first, and so on, in descending order. By default, the original source volume is selected.
File path text box	Enables you to type the file path to which you want data restored. The path you enter must already exist.
Name	Displays the names of the available folders for the selected volume. Clicking a folder in the Name list displays the subfolders, if any exist. Files contained in the folders are not displayed. Clicking after you have selected a folder moves you up one level in the directory structure.

Command buttons

The command buttons enable you to perform the following tasks:

Select Directory	Applies your selections and closes the Browse Directories dialog box. If no directory is selected, this button is disabled.
Cancel	Discards your selections and closes the Browse Directories dialog box.

Related references

Restore dialog box on page 263

Job details page

The Job details page enables you to view status and other information about specific protection job tasks that are running, that are queued, or that have completed. You can use this information to monitor protection job progress and to troubleshoot job failures.

Job summary

The job summary displays the following information:

- Job ID
- Type
- State
- Submitted Time
- Completed Time
- Duration

Command buttons

The command buttons enable you to perform the following tasks:

Refresh	Refreshes the task list and the properties associated with each task.
View Jobs	Returns you to the Jobs page.

Job tasks list

The Job tasks list displays in a table all the tasks associated with a specific job and the properties related to each task.

Started Time	Displays the day and time the task started. By default, the most recent tasks are displayed at the top of the column and older tasks are displayed at the bottom.
Type	Displays the type of task.
State	The state of a particular task: **Completed** The task has finished. **Queued** The task is about to run. **Running** The task is running. **Waiting** A job has been submitted and some associated tasks are waiting to be queued and executed.
Status	Displays the task status:

Error ( )	The task failed.
Normal ( )	The task succeeded.
Skipped ( )	A task failed, resulting in subsequent tasks being skipped.

Duration	Displays the elapsed time since the task began.
Completed Time	Displays the time the task completed. By default, this column is hidden.
Task ID	Displays the GUID that identifies an individual task for a job. The column can be sorted and filtered. By default, this column is hidden.
Dependency order	Displays an integer representing the sequence of tasks in a graph, with zero being assigned to the first task. By default, this column is hidden.
Task Details pane	Displays additional information about each job task, including the task name, task description, and, if the task failed, a reason for the failure.
Task Messages pane	Displays messages specific to the selected task. Messages might include a reason for the error and suggestions for resolving it. Not all tasks display task messages.

Related tasks

Viewing job details on page 249

Relationships Thresholds Setup Options dialog box

The Relationships Thresholds Setup Options dialog box enables you to configure global lag warning and error threshold values for unmanaged protection relationships so that you are notified and can take action when lag or threshold errors occur. Changes to these settings are applied during the next scheduled update.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Events are generated when a threshold is breached. You can take corrective actions for such events. Lag threshold settings for unmanaged relationships are enabled by default.

- *Lag Thresholds for Unmanaged Relationships* on page 267
- *Command buttons* on page 268

Lag Thresholds for Unmanaged Relationships area

The Lag area enables you set unmanaged relationship lag thresholds for the following conditions:

Warning	Specifies the percentage at which the lag duration equals or exceeds the lag warning threshold:

- Default value: 150 percent
- Events generated: SnapMirror Relationship Lag Warning or SnapVault Relationship Lag Warning
- Event severity: Warning

Error Specifies the percentage at which the lag duration equals or exceeds the lag error threshold:

- Default value: 250 percent
- Events generated: SnapMirror Relationship Lag Error or SnapVault Relationship Lag Error
- Event severity: Error

Command buttons

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

Restore to Factory Defaults	Enables you to restore the configuration settings to the factory default values.
Save	Saves the configuration settings for the selected option.
Save and Close	Saves the configuration settings for the selected option and closes the Setup Options dialog box.
Cancel	Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Editing unmanaged relationship lag threshold settings on page 99

Edit Aggregate Thresholds dialog box

You can configure alerts to send notifications when an event related to an aggregate's capacity is generated, and you can take corrective actions for the event. For example, for the Aggregate Full threshold, you can configure an alert to send notification when the condition persists over a specified period.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

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

- *Capacity* on page 269
- *Growth* on page 269
- *Snapshot Copies* on page 269
- *Command buttons* on page 270

Capacity area

The Capacity area enables you to set the following aggregate capacity threshold conditions:

Space Nearly Full	Specifies the percentage at which an aggregate is considered to be nearly full. It also displays the size of the aggregate corresponding to the specified threshold value. You can also use the slider to set the threshold value.
Space Full	Specifies the percentage at which an aggregate is considered full. It also displays the size of the aggregate corresponding to the specified threshold value. You can also use the slider to set the threshold value.
Nearly Overcommitted	Specifies the percentage at which an aggregate is considered to be nearly overcommitted.
Overcommitted	Specifies the percentage at which an aggregate is considered to be overcommitted.
Days Until Full	Specifies the number of days remaining before the aggregate reaches full capacity.

Growth

The Growth area enables you to set the following threshold condition for aggregate growth:

Growth Rate	Specifies the percentage at which an aggregate's growth rate is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event.
Growth Rate Sensitivity	Specifies the factor that is applied to the standard deviation of an aggregate's growth rate. If the growth rate exceeds the factored standard deviation, an Aggregate Growth Rate Abnormal event is generated. A lower value for growth rate sensitivity indicates that the aggregate is highly sensitive to changes in the growth rate. **Note:** If you modify the growth rate sensitivity for aggregates at the global threshold level, the change is also applied to the growth rate sensitivity for volumes at the global threshold level.

Snapshot copies area

The Snapshot copies area enables you to set the following Snapshot reserve threshold conditions:

Snapshot Reserve Full	Specifies the percentage at which an aggregate has consumed all its space reserved for Snapshot copies. You can also use the slider to set the threshold value.

Command buttons

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

Restore to Global Defaults	Enables you to restore the threshold settings to the current values that are set at the global level.
Save	Saves all the threshold settings.
Save and Close	Saves all the threshold settings and then closes the Edit Aggregate Thresholds dialog box.
Cancel	Ignores the changes (if any) to the threshold settings and closes the Edit Aggregate Thresholds dialog box.

Related concepts

Understanding capacity events and thresholds for node root aggregates on page 119

Related tasks

Editing aggregate threshold settings on page 101

Edit Volume Thresholds dialog box

You can configure alerts to send notifications when an event related to a volume's capacity is generated, and you can take corrective actions for the event. For example, for the Volume Full threshold, you can configure an alert to send notification when the condition persists over a specified period.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

- *Capacity* on page 270
- *Snapshot Copies* on page 271
- *Quota* on page 271
- *Growth* on page 271
- *Inodes* on page 272
- *Command buttons* on page 272

Capacity area

The Capacity area enables you to set the following volume capacity threshold conditions:

Space Nearly Full	Specifies the percentage at which a volume is considered to be nearly full. It also displays the size of the volume corresponding to the specified threshold value. You can also use the slider to set the threshold value.

Space Full	Specifies the percentage at which a volume is considered full. It also displays the size of the volume corresponding to the specified threshold value. You can also use the slider to set the threshold value.
Days Until Full	Specifies the number of days remaining before the volume reaches full capacity.

Snapshot Copies

The Snapshot Copies area enables you to set the following threshold conditions for the Snapshot copies in the volume.

Snapshot Reserve Full	Specifies the percentage at which the space reserved for Snapshot copies is considered full.
Days Until Full	Specifies the number of days remaining before the space reserved for Snapshot copies reaches full capacity.
Count	Specifies the number of Snapshot copies that can be created on a volume before the system generates the Too Many Snapshot Copies event. **Note:** This field is applicable only for volumes in a cluster running Data ONTAP 8.2 or later.

Quota area

The Quota area enables you to set the following volume quota threshold conditions:

Nearly Overcommitted	Specifies the percentage at which a volume is considered to be nearly overcommitted by qtree quotas.
Overcommitted	Specifies the percentage at which a volume is considered to be overcommitted by qtree quotas.

Growth area

The Growth area enables you to set the following threshold condition for volume growth:

Growth Rate	Specifies the percentage at which a volume's growth rate is considered to be normal before the system generates a Volume Growth Rate Abnormal event.
Growth Rate Sensitivity	Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated. A lower value for growth rate sensitivity indicates that the volume is highly sensitive to changes in the growth rate.

Note: If you modify the growth rate sensitivity for volumes at the global threshold level, the change is also applied to the growth rate sensitivity for aggregates at the global threshold level.

Inodes

The Inodes enables you to set the following threshold conditions for inodes:

Nearly Full	Specifies the percentage at which a volume is considered to have consumed most of its inodes. You can also use the sliders to set the threshold value.
Full	Specifies the percentage at which a volume is considered to have consumed all of its inodes. You can also use the sliders to set the threshold value.

Command buttons

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

Restore to Global Defaults	Enables you to restore the threshold settings to the current values that are set at the global level.
Save	Saves all the threshold settings.
Save and Close	Saves all the threshold settings and then closes the Aggregates dialog box.
Cancel	Ignores the changes (if any) to the threshold settings and closes the Aggregates dialog box.

Related tasks

Editing volume threshold settings on page 102

Configure Protection dialog box

You can use the Configure Protection dialog box to create SnapMirror and SnapVault relationships for all read, write, and data protection volumes on clusters running Data ONTAP 8.2 or later to ensure data on a source or primary volume is replicated.

- *Source tab* on page 273
- *SnapMirror tab* on page 273
- *SnapVault tab* on page 275
- *Command buttons* on page 277

Source tab

Topology view

Displays a visual representation of the relationship you are creating. The source in the topology is highlighted by default.

Source Information

Displays details about the selected source volumes, including the following information:

- Source cluster name
- Source SVM name
- Cumulative volume total size
  Displays the total size of all the source volumes selected.
- Cumulative volume used size
  Displays the cumulative volume used size for all the selected source volumes.
- Source volume
  Displays in a table the following information:
  - Source Volume
    Displays the names of the selected source volumes.
  - Type
    Displays the volume type.
  - Snapshot Copy
    Displays the Snapshot copy that is used for the baseline transfer. If the source volume is read/write, the value of Default in the Snapshot copy column indicates that a new Snapshot copy is created by default and is used for the baseline transfer; however, if the volume is a data protection volume, the value of Default in the Snapshot copy column indicates that no new Snapshot copy is created, and all existing Snapshot copies are transferred to the destination. Clicking the Snapshot copy value displays a list of Snapshot copies from which you can select an existing Snapshot copy to use for the baseline transfer. You cannot select a different default Snapshot copy if the source type is data protection.

SnapMirror tab

When you create a SnapMirror relationship, the SnapMirror tab enables you to specify a destination cluster, Storage Virtual Machine (SVM), and aggregate for a protection relationship, as well as naming convention for destinations. You can also specify a SnapMirror policy and schedule.

Topology view

Displays a visual representation of the relationship you are creating. The SnapMirror destination resource in the topology is highlighted by default.

Destination Information

Enables you to select the destination resources for a protection relationship:

- Advanced link

Launches the Advanced Destination Settings dialog box when you are creating a SnapMirror relationship.

- Cluster

  Lists the clusters available as the protection destination hosts. This field is required.
- Storage Virtual Machine (SVM)

  Lists the SVMs available on the selected cluster. A cluster must be selected before the SVM list is populated. This field is required.
- Aggregate

  Lists the aggregates available on the selected SVM. A cluster must be selected before the aggregate list is populated. This field is required. The aggregate list displays the following information:

  - Rank

    When multiple aggregates satisfy all the requirements for a destination, the rank indicates the priority in which the aggregate is listed according to the following conditions:

    1. An aggregate located on a different node than the source volume node is preferred to enable fault domain separation.
    2. An aggregate on a node with fewer volumes is preferred to enable load balancing across nodes in a cluster.
    3. An aggregate that has more free space than other aggregates is preferred to enable capacity balancing.

    A rank of 1 means the aggregate is the most preferred according to the three criteria.
  - Aggregate Name

    The name of the aggregate.
  - Available Capacity
  - The amount of space available on the aggregate for data.
  - Resource Pool

    The name of the resource pool to which the aggregate belongs.
- Naming Convention

  Specifies the default naming convention that is applied to the destination volume. You can accept the naming convention provided, or you can create a new one. The naming convention can have the following attributes: %C, %M, %V, %N, where %C is the cluster name, %M is the SVM name, %V is the source volume, and %N is the topology destination node name.

  The naming convention field is highlighted in red if your entry is invalid. Clicking the “Preview Name” link displays a preview of the naming convention you entered, and the preview text updates dynamically as you

type a naming convention in the text field. A suffix between 001 and 999 is appended to the destination name when the relationship is created, replacing the *nnn* that displays in the preview text, with 001 being assigned first, 002 assigned second, and so on.

Relationship Settings

Enables you to specify the maximum transfer rate, SnapMirror policy, and schedule that the protection relationship uses:

- Max Transfer Rate

  Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running Data ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored.
- SnapMirror Policy

  Specifies the Data ONTAP SnapMirror policy for the relationship. The default is DPDefault.
- Create Policy

  Launches the Create SnapMirror Policy dialog box dialog box, which enables you to create and use a new SnapMirror policy.
- SnapMirror Schedule

  Specifies the Data ONTAP SnapMirror policy for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.
- Create Schedule

  Launches the Create Schedule dialog box, which enables you to create a new SnapMirror schedule.

SnapVault tab

When you create a SnapVault relationship, the SnapVault tab enables you to specify a secondary cluster, Storage Virtual Machine (SVM), and aggregate for a protection relationship, as well as the naming convention for secondary volumes. You can also specify a SnapVault policy and schedule.

Topology view

Displays a visual representation of the relationship you are creating. The SnapVault secondary resource in the topology is highlighted by default.

Secondary Information

Enables you to select the secondary resources for a protection relationship:

- Advanced link

  Launches the Advanced Secondary Settings dialog box.
- Cluster

Lists the clusters available as the secondary protection hosts. This field is required.

- Storage Virtual Machine (SVM)

  Lists the SVMs available on the selected cluster. A cluster must be selected before the SVM list is populated. This field is required.
- Aggregate

  Lists the aggregates available on the selected SVM. A cluster must be selected before the aggregate list is populated. This field is required. The aggregate list displays the following information:

  - Rank

    When multiple aggregates satisfy all the requirements for a destination, the rank indicates the priority in which the aggregate is listed according to the following conditions:

    1. An aggregate located on a different node than the primary volume node is preferred to enable fault domain separation.
    2. An aggregate on a node with fewer volumes is preferred to enable load balancing across nodes in a cluster.
    3. An aggregate that has more free space than other aggregates is preferred to enable capacity balancing.

    A rank of 1 means the aggregate is the most preferred according to the three criteria.
  - Aggregate Name

    The name of the aggregate.
  - Available Capacity
  - The amount of space available on the aggregate for data.
  - Resource Pool

    The name of the resource pool to which the aggregate belongs.
- Naming Convention

  Specifies the default naming convention that is applied to the secondary volume. You can accept the naming convention provided, or you can create a new one. The naming convention can have the following attributes: %C, %M, %V, %N, where %C is the cluster name, %M is the SVM name, %V is the source volume, and %N is the topology secondary node name.

  The naming convention field is highlighted in red if your entry is invalid. Clicking the “Preview Name” link displays a preview of the naming convention you entered, and the preview text updates dynamically as you type a naming convention in the text field. If you type an invalid value, the invalid information displays as red question marks in the preview area.. A suffix between 001 and 999 is appended to the secondary name

when the relationship is created, replacing the *nnn* that displays in the preview text, with 001 being assigned first, 002 assigned second, and so on.

Relationship Settings

Enables you to specify the maximum transfer rate, SnapVault policy, and SnapVault schedule that the protection relationship uses:

- Max Transfer Rate

  Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running Data ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored.
- SnapVault Policy

  Specifies the Data ONTAP SnapVault policy for the relationship. The default is XDPDefault.
- Create Policy

  Launches the Create SnapVault Policy dialog box, which enables you to create and use a new SnapVault policy.
- SnapVault Schedule

  Specifies the Data ONTAP SnapVault schedule for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.
- Create Schedule

  Launches the Create Schedule dialog box, which enables you to create a SnapVault schedule.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel	Discards the selections and closes the Configure Protection dialog box.
Apply	Applies your selections and begins the protection process.

Related tasks

Creating a SnapVault protection relationship from the Volumes page on page 216
Creating a SnapVault protection relationship from the Volume details page on page 217
Creating a SnapMirror protection relationship from the Volumes page on page 219
Creating a SnapMirror protection relationship from the Volume details page on page 218

Related references

Advanced Secondary Settings dialog box on page 261
Advanced Destination Settings dialog box on page 262

Create Schedule dialog box

The Create Schedule dialog box enables you to create a basic or advanced protection schedule for SnapMirror and SnapVault relationship transfers. You might create a new schedule to increase the frequency of data transfers due to frequent data updates, or you might create a less frequent schedule when data changes infrequently.

Destination Cluster

The name of the cluster you selected in the SnapVault tab or SnapMirror tab of the Configure Protection dialog box.

Schedule Name

The name you provide for the schedule. Schedule names can consist of the characters A through Z, a through z, 0 through 9, as well as any of the following special characters: ! @ # $ % ^ & * ( ) _ -. Schedule names may not include the following characters: < >.

Basic or Advanced

The schedule mode you want to use.

Basic mode includes the following elements:

- Repeat
  How often a scheduled transfer occurs. Choices include hourly, daily, and weekly.
- Day
  When a repeat of weekly is selected, the day of the week a transfer occurs.
- Time
  When Daily or Weekly is selected, the time of day a transfer occurs.

Advanced mode includes the following elements:

- Months
  A comma-separated numerical list representing the months of the year. Valid values are 0 through 11, with zero representing January, and so on. This element is optional. Leaving the field blank implies that transfers occur every month.
- Days
  A comma-separated numerical list representing the day of the month. Valid values are 1 through 31. This element is optional. Leaving the field blank implies that a transfer occurs every day of the month.
- Weekdays
  A comma-separated numerical list representing the days of the week. Valid values are 0 through 6, with 0 representing Sunday, and so on. This element is optional. Leaving the field blank implies that a transfer occurs every day of the

week. If a day of the week is specified but a day of the month is not specified, a transfer occurs only on the specified day of the week and not every day.

- Hours

  A comma-separated numerical list representing the number of hours in a day. Valid values are 0 through 23, with 0 representing midnight. This element is optional.
- Minutes

  A comma-separated numerical list representing the minutes in an hour. Valid values are 0 through 59. This element is required.

Related tasks

Creating SnapMirror and SnapVault schedules on page 221

Create SnapMirror Policy dialog box

The Create SnapMirror Policy dialog box enables you to create a policy to set the priority for SnapMirror transfers. You use policies to maximize the efficiency of transfers from the source to the destination.

Destination Cluster	The name of the cluster you selected in the SnapMirror tab of the Configure Protection dialog box.
Destination Storage Virtual Machine (SVM)	The name of the SVM you selected in the SnapMirror tab of the Configure Protection dialog box.
Policy Name	The name you provide for the new policy. Policy names can consist of the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underscore (_).
Transfer Priority	The priority at which the transfer is run. You can select either Normal or Low. Transfer relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority.
Comment	An optional field in which you can add comments about the policy.
Transfer Restart	Indicates what restart action to take when a transfer is interrupted by an abort operation or any type of failure, such as a network outage. You can select one of the following: • Always Specifies that a new Snapshot copy is created before restarting a transfer, then, if one exists, the transfer is restarted from a checkpoint, followed by an incremental transfer from the newly created Snapshot copy.. • Never Specifies that interrupted transfers are never restarted.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel	Discards the selections and closes the Configure Protection dialog box.
Apply	Applies your selections and begins the protection process.

Related tasks

Creating SnapMirror policies on page 224

Create SnapVault Policy dialog box

The Create SnapVault Policy dialog box enables you to create a policy to set the priority for SnapVault transfers. You use policies to maximize the efficiency of transfers from the primary to the secondary volume.

Destination Cluster	The name of the cluster that you selected in the SnapVault tab of the Configure Protection dialog box.
Destination Storage Virtual Machine (SVM)	The name of the SVM that you selected in the SnapVault tab of the Configure Protection dialog box.
Policy Name	The name you provide for the new policy. Policy names can consist of the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underscore (_).
Transfer Priority	The priority at which the transfer is run. You can select either Normal or Low. Transfer relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority. The default setting is Normal.
Comment	An optional field in which you can a add comment of up to 255 characters about the SnapVault policy.
Ignore Access Time	Specifies whether incremental transfers are ignored for files that have only their access time changed.

Replication Label	Lists in a table the rules associated with Snapshot copies selected by Data ONTAP that have a specific replication label in a policy. The following information and actions are also available:
	• Command buttons
	The command buttons enable you to perform the following actions:
	  • Add
	  Enables you to create a Snapshot copy label and retention count.
	  • Edit Retention Count

Enables you to change the retention count for an existing Snapshot copy label. The retention count must be a number between 1 and 251. The sum of all retention counts for all rules cannot exceed 251.

- Delete

  Enables you to delete an existing Snapshot copy label.

- Snapshot Copy Label

  Displays the Snapshot copy label. If you select one or more volumes with the same local Snapshot copy policy, an entry for each label in the policy is displayed. If you select multiple volumes that have two or more local Snapshot copy policies, the table displays all labels from all policies

- Schedule

  Displays the schedule associated with each Snapshot copy label. If a label has more than one schedule associated with it, the schedules for that label are displayed in a comma-separated list. If you select multiple volumes with the same label but with different schedules, the schedule displays “Various” to indicate that more than one schedule is associated with the selected volumes.

- Destination Retention Count

  Displays the number of Snapshot copies with the specified label that are retained on the SnapVault secondary. Retention counts for labels with multiple schedules displays the sum of retention counts of each label and schedule pair. If you select multiple volumes with two or more local Snapshot copy policies, the retention count is empty.

Related tasks

Creating SnapVault policies on page 224

Edit Relationship dialog box

You can edit an existing protection relationship to change the maximum transfer rate, the protection policy, or the protection schedule.

Destination Information

Destination Cluster	The name of the selected destination cluster.
Destination SVM	The name of the selected Storage Virtual Machine (SVM)
Relationship Settings	Enables you to specify the maximum transfer rate, SnapMirror policy, and schedule that the protection relationship uses: • Max Transfer Rate Specifies the maximum rate at which baseline data is transferred between clusters over the network. When selected, network bandwidth is limited

to the value you specify. You can enter a numerical value and then select either kilobytes per second (KBps), megabytes per second (MBps), gigabytes per second (GBps), or terabytes per second (TBps). If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. If the primary cluster and the secondary cluster are the same, this setting is disabled.

- SnapMirror Policy

  Specifies the Data ONTAP SnapMirror policy for the relationship. The default is DPDefault.
- Create Policy

  Launches the Create SnapMirror Policy dialog box, which enables you to create and use a new SnapMirror policy.
- SnapMirror Schedule

  Specifies the Data ONTAP SnapMirror policy for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.
- Create Schedule

  Launches the Create Schedule dialog box, which enables you to create a new SnapMirror schedule.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel	Discards the selections and closes the Configure Protection dialog box.
Submit	Applies your selections and closes the Edit Relationship dialog box.

Related tasks

Editing protection relationships from the Volumes page on page 222
Editing protection relationships from the Volume details page on page 223

Initialize/Update dialog box

The Initialize/Update dialog box enables you to perform a first-time baseline transfer on a new protection relationship, or to update a relationship if it is already initialized and you want to perform a manual, unscheduled, incremental update.

Transfer Options tab

The Transfer Options tab enables you to change the initialization priority of a transfer and to change the bandwidth used during transfers.

Transfer Priority The priority at which the transfer is run. You can select either Normal or Low. Relationships with policies that specify a normal transfer priority run before those that specify a low transfer priority. Normal is selected by default.

Max Transfer Rate Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running Data ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored. If you select more than one relationship with different maximum transfer rates, you can specify one of the following maximum transfer rate settings:

- Use values specified during individual relationship setup or edit
  When selected, initialization and update operations use the maximum transfer rate specified at the time of each relationship's creation or edit. This field is available only when multiple relationships with different transfer rates are being initialized or updated.
- Unlimited
  Indicates that there is no bandwidth limitation on transfers between relationships. This field is available only when multiple relationships with different transfer rates are being initialized or updated.
- Limit bandwidth to
  When selected, network bandwidth is limited to the value you specify. You can enter a numerical value and then select either kilobytes per second (KBps), Megabytes per second (MBps), Gigabytes per second (GBps), or Terabytes per second (TBps).

Source Snapshot Copies tab

The Source Snapshot Copies tab displays the following information about the source Snapshot copy that is used for the baseline transfer:

Source Volume Displays the names of the corresponding source volumes.

Destination Volume Displays the names of the selected destination volumes.

Source Type Displays the volume type. The type can be either Read/write or Data Protection.

Snapshot Copy Displays the Snapshot copy that is used for the data transfer. Clicking the Snapshot copy value displays the Select Source Snapshot Copy dialog box, in which can select a specific Snapshot copy for your transfer, depending on the type of protection relationship that you have and the operation that you are performing. The option to specify a difference Snapshot copy is not available for data protection type sources.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel	Discards your selections and closes the Initialize/Update dialog box.
Submit	Saves your selections and starts the initialize or update job.

Related tasks

Initializing or updating protection relationships from the Volume details page on page 231
Initializing or updating protection relationships from the Volumes page on page 232

Related references

Select Source Snapshot Copy dialog box on page 285

Resynchronize dialog box

The Resynchronize dialog box enables you to resynchronize data on a SnapMirror or SnapVault relationship that was previously broken and then the destination was made a read/write volume. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Resynchronization Options tab

The Resynchronization Options tab enables you to set the transfer priority and the maximum transfer rate for the protection relationship that you are resynchronizing.

Transfer Priority	The priority at which the transfer is run. You can select either Normal or Low. Relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority.
Max Transfer Rate	Specifies the maximum rate at which data is transferred between clusters over the network. When selected, network bandwidth is limited to the value that you specify. You can enter a numerical value and then select either kilobytes per second (KBps), megabytes per second (MBps), gigabytes per second (GBps), or TBps. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running Data ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is disabled.

Source Snapshot Copies tab

The Source Snapshot Copies tab displays the following information about the source Snapshot copy that is used for the baseline transfer:

Source Volume	Displays the names of the corresponding source volumes.

Destination Volume	Displays the names of the selected destination volumes.
Source Type	Displays the volume type: either Read/write or Data Protection.
Snapshot Copy	Displays the Snapshot copy that is used for the data transfer. Clicking the Snapshot copy value displays the Select Source Snapshot Copy dialog box, in which can select a specific Snapshot copy for your transfer, depending on the type of protection relationship you have and the operation you are performing.

Command buttons

Submit	Begins the resynchronization process and closes the Resynchronize dialog box.
Cancel	Cancels your selections and closes the Resynchronize dialog box.

Related tasks

Resynchronizing protection relationships from the Volume details page on page 233
Resynchronizing protection relationships from the Volumes page on page 234

Related references

Select Source Snapshot Copy dialog box on page 285

Select Source Snapshot Copy dialog box

You use the Select Source Snapshot Copy dialog box to select a specific Snapshot copy to transfer data between protection relationships, or you select the default behavior, which varies depending on whether you are initializing, updating, or resynchronizing a relationship, and whether the relationship is a SnapMirror or SnapVault.

Default

Enables you to select the default behavior for determining which Snapshot copy is used for initialize, update, and resynchronize transfers for SnapVault and SnapMirror relationships.

If you are performing a SnapVault transfer, the default behavior for each operation is as follows:

Operation	**Default SnapVault behavior when source is read/write**	**Default SnapVault behavior when source is Data Protection (DP)**
Initialize	Creates a new Snapshot copy and transfers it.	Transfers the last exported Snapshot copy.

Operation	**Default SnapVault behavior when source is read/write**	**Default SnapVault behavior when source is Data Protection (DP)**
Update	Transfers only labeled Snapshot copies, as specified in the policy.	Transfers the last exported Snapshot copy.
Resynchronize	Transfers all labeled Snapshot copies created after the newest common Snapshot copy.	Transfers the newest labeled Snapshot copy.

If you are performing a SnapMirror transfer, the default behavior for each operation is as follows:

Operation	**Default SnapMirror behavior**	**Default SnapMirror behavior when relationship is second hop in a SnapMirror to SnapMirror cascade**
Initialize	Creates a new Snapshot copy and transfers it and all Snapshot copies created prior to the new Snapshot copy.	Transfers all Snapshot copies from the source.
Update	Creates a new Snapshot copy and transfers it and all Snapshot copies created prior to the new Snapshot copy.	Transfers all Snapshot copies.
Resynchronize	Creates a new Snapshot copy and then transfers all Snapshot copies from the source.	Transfers all Snapshot copies from the secondary volume to the tertiary volume, and deletes any data added after creation of the newest common Snapshot copy.

Existing Snapshot Copy

Enables you to select an existing Snapshot copy from the list if Snapshot copy selection is allowed for that operation.

Snapshot Copy Displays the existing Snapshot copies from which you can select for a transfer.

Date Created Displays the date and time the Snapshot copy was created. Snapshot copies are listed from most recent to least recent, with the most recent at the top of the list.

If you are performing a SnapVault transfer and you want to select an existing Snapshot copy to transfer from a source to a destination, the behavior for each operation is as follows:

Operation	**SnapVault behavior when specifying a Snapshot copy**	**SnapVault behavior when specifying a Snapshot copy in a cascade**
Initialize	Transfers the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Update	Transfers the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Resynchronize	Transfers the selected Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

If you are performing a SnapMirror transfer and you want to select an existing Snapshot copy to transfer from a source to a destination, the behavior for each operation is as follows:

Operation	**SnapMirror behavior when specifying a Snapshot copy**	**SnapMirror behavior when specifying a Snapshot copy in a cascade**
Initialize	Transfers all Snapshot copies on the source, up to the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Update	Transfers all Snapshot copies on the source, up to the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Resynchronize	Transfers all Snapshot copies from the source, up to the selected Snapshot copy, and then deletes any data added after creation of the newest common Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

Command buttons

The command buttons enable you to perform the following tasks:

Submit Submits your selections and closes the Select Source Snapshot Copy dialog box.

Cancel Discards your selections and closes the Select Source Snapshot Copy dialog box.

Related tasks

Resynchronizing protection relationships from the Volume details page on page 233
Resynchronizing protection relationships from the Volumes page on page 234

Related references

Resynchronize dialog box on page 284
Initialize/Update dialog box on page 282

Reverse Resync dialog box

When you have a protection relationship that is broken because the source volume is disabled and the destination is made a read/write volume, reverse resynchronization enables you to reverse the direction of the relationship so that the destination becomes the new source and the source becomes the new destination.

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write, while you repair or replace the source, update the source, and reestablish the relationship. When you perform a reverse resync operation, data on the source that is newer than the data on the common Snapshot copy is deleted.

Before reverse resync

Displays the source and destination of a relationship before a reverse resync operation.

Source Volume	The name and location of the source volume before a reverse resync operation .
Destination Volume	The name and location of the destination volume before a reverse resync operation .

After reverse resync

Displays what the source and destination of a relationship is after a reserve resync operation.

Source Volume	The name and location of the source volume after a reverse resync operation.
Destination Volume	The name and location of the destination volume after a reverse resync operation.

Command buttons

The command buttons enable you to perform the following actions:

Submit	Begins the reverse resynchronization process.
Cancel	Closes the Reverse Resync dialog box without initiating a reverse resync operation.

Related tasks

Reversing protection relationships from the Volume details page on page 235
Reversing protection relationships from the Volumes page on page 236

Executing protection workflows using OnCommand Workflow Automation

You can integrate OnCommand Workflow Automation with Unified Manager to execute workflows for your storage classes and monitor SVMs with Infinite Volume that do not have storage classes.

Pairing OnCommand Workflow Automation with Unified Manager

You can integrate Workflow Automation with Unified Manager over a secure connection. This enables protection features such as SnapMirror and SnapVault configuration workflows, as well as commands for managing SnapMirror relationships.

Before you begin

The following information must be available:

- Name of a database user in Unified Manager
- Host address, port number 443, user name, and password for the OnCommand Workflow Automation setup

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

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

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Add-ons > Workflow Automation**.
3. In the **Set Up OnCommand Workflow Automation** dialog box, type the host name or IP address and the user name and password.

   You must use Unified Manager server port 443.
4. Click **Save and Close**.
5. If you use a self-signed certificate, click **Yes** to authorize the security certificate.

   The Workflow Automation Options Changed dialog box displays.
6. Click **Yes** to reload the web UI and add the Workflow Automation features.

Related references

Set Up OnCommand Workflow Automation dialog box on page 292

Removing OnCommand Workflow Automation setup from Unified Manager

You can remove the OnCommand Workflow Automation setup from Unified Manager when you no long want to use OnCommand Workflow Automation.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Add-ons > Workflow Automation**.
3. In the **Set Up OnCommand Workflow Automation** dialog box, click **Remove Setup**.

Related references

Set Up OnCommand Workflow Automation dialog box on page 292

What happens when OnCommand Workflow Automation is reinstalled or upgraded

When you reinstall or upgrade OnCommand Workflow Automation (WFA), the WFA SSL certificate changes and the Unified Manager and WFA integration is interrupted, therefore preventing the execution of the protection workflows.

After you reinstall or upgrade OnCommand Workflow Automation, you must execute the Unified Manager and WFA integration process on the Unified Manager server to update the settings. To execute the integration process, you must provide the OnCommand Workflow Automation credentials along with the DB user password.

Description of OnCommand Workflow Automation setup windows and dialog boxes

You can set up OnCommand Workflow Automation in Unified Manager by using the Set Up OnCommand Workflow Automation dialog box.

Set Up OnCommand Workflow Automation dialog box

The Set Up OnCommand Workflow Automation dialog box enables you to configure the settings to integrate OnCommand Workflow Automation with Unified Manager. You can also add, modify, or delete the settings.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Unified Manager Database User

This area enables you to enter the credentials of a database user that is required for pairing Unified Manager with OnCommand Workflow Automation:

Name Enables you to specify the user name of a database user that can be used to access data in the Unified Manager database. By default, no database user is selected. You can select a database user from the drop-down list.

Password Enables you to specify a password for the specified user name.

OnCommand Workflow Automation Credentials

This area enables you to enter the credentials of an OnCommand Workflow Automation account that is required for pairing with Unified Manager:

Host Specifies the address of the OnCommand Workflow Automation host server, which is used to pair with Unified Manager.

Port Specifies the port number of the OnCommand Workflow Automation host server. If the port number is not specified, the default value of the port is 443.

Username Enables you to specify a user name that can be used to log in to OnCommand Workflow Automation.

Password Enables you to specify a password for the specified user name.

Command buttons

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

Remove Setup	Removes the OnCommand Workflow Automation setup from Unified Manager.
Save	Saves the configuration settings for the selected option.
Save and Close	Saves the configuration settings for the selected option and closes the Setup Options dialog box.
Cancel	Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Pairing OnCommand Workflow Automation with Unified Manager on page 290
Removing OnCommand Workflow Automation setup from Unified Manager on page 291

Managing and monitoring Infinite Volumes

You can monitor the capacity and availability of your SVMs with Infinite Volume. You can manage the content placement in your SVM with Infinite Volume by creating rules and data policy.

Viewing the details of SVMs with Infinite Volume

You can use the Storage Virtual Machines page to view detailed information about SVMs with Infinite Volume that are monitored by Unified Manager. You can view details such as the capacity, configuration, and data policy and rules associated with the Infinite Volume.

Before you begin

- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.
- The cluster containing the SVM with Infinite Volume must be added to the Unified Manager database.

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. In the Storage Virtual Machines page, use the column filter in **Allowed Volume Type** to list the Infinite Volumes that are monitored.
3. View the complete details of the SVM with Infinite Volume by clicking the SVM name.

Related tasks

Adding a user on page 316
Adding an alert on page 77
Editing the Infinite Volume threshold settings on page 295
Editing the threshold settings of storage classes on page 296
Viewing rules on page 301

Viewing the constituents of an Infinite Volume

You can use the Volumes page to view the list of constituents in your Infinite Volume. You can view details such as the constituent state, the SVM with Infinite Volume that contains the constituent,

junction path of the constituent, aggregate that contains the constituent, as well as the available, used, and total data capacity of the constituent.

Before you begin

The following requirements must be met:

- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.
- The cluster containing the SVM with Infinite Volume must be added to the Unified Manager database.

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. Click the name of the appropriate SVM with Infinite Volume.
3. In the **Storage Virtual Machine details** page, click **Volumes** in the **Related Devices** pane.

   The list of constituents is displayed in the Volumes page.

Related concepts

What a namespace constituent is on page 298
What data constituents are on page 298
What a namespace mirror constituent is on page 298

Related tasks

Adding a user on page 316

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 be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the required SVM with Infinite Volume.

3. In the **Storage Virtual Machine details** page, click **Actions > Edit Thresholds**.
4. In the **Edit SVM with Infinite Volume Thresholds** dialog box, modify the thresholds as required.
5. Click **Save and Close**.

Related tasks

Adding a user on page 316

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 be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the required SVM with Infinite Volume.
3. In the **Storage Virtual Machine details** page, click **Actions > Edit Thresholds**.
4. In the **Edit Storage Class Thresholds** dialog box, modify the thresholds as required.
5. Click **Save and Close**.

Related references

Unified Manager roles and capabilities on page 322

Understanding Infinite Volumes

An Infinite Volume is a logical storage unit that you can use to provide a large, scalable data container with a single namespace and a single mount point. Understanding some of the basic concepts of Infinite Volumes helps you to monitor and manage your SVMs with Infinite Volume.

What an Infinite Volume is

An Infinite Volume is a single, scalable volume that can store up to 2 billion files and tens of petabytes of data.

With an Infinite Volume, you can manage multiple petabytes of data in one large logical entity and clients can retrieve multiple petabytes of data from a single junction path for the entire volume.

An Infinite Volume uses storage from multiple aggregates on multiple nodes. You can start with a small Infinite Volume and expand it nondisruptively by adding more disks to its aggregates or by providing it with more aggregates to use.

Related concepts

Maximum number of files an Infinite Volume can store on page 297

Maximum number of files an Infinite Volume can store

In most cases, an Infinite Volume can hold up to 2 billion files. If an Infinite Volume is relatively small, its maximum number of files might be less than 2 billion.

The maximum number of files that an Infinite Volume can hold is determined by the size of its namespace constituent. If the namespace constituent is 10 TB, the Infinite Volume can hold 2 billion files. If the namespace constituent is less than 10 TB, the Infinite Volume can hold proportionally fewer files.

The size of the namespace constituent is roughly proportional to the size of the Infinite Volume, depending on several factors, such as the namespace constituent's 10 TB maximum size, the available space in the aggregate that holds the namespace constituent, and the SnapDiff setting.

For a two-node Infinite Volume or a multi-node Infinite Volume without SnapDiff enabled, setting the Infinite Volume to a size of 80 TB or greater typically creates a namespace constituent of 10 TB.

The file count not only includes regular files, but also other file system structures, such as directories and symbolic links.

Related concepts

What an Infinite Volume is on page 297

What a storage class is

A storage class is a definition of aggregate characteristics and volume settings. You can define different storage classes and associate one or more storage classes with an Infinite Volume. You must use OnCommand Workflow Automation to define workflows for your storage class requirements and to assign storage classes to Infinite Volumes.

You can define the following characteristics for a storage class:

- Aggregate characteristics, such as the type of disks to use

- Volume settings, such as compression, deduplication, and volume guarantee

For example, you can define a storage class that uses only aggregates with SAS disks and the following volume settings: thin provisioning with compression and deduplication enabled.

What a namespace constituent is

Each Infinite Volume has a single namespace constituent that maps directory information and file names to the file's physical data location within the Infinite Volume.

Clients are not aware of the namespace constituent and do not interact directly with it. The namespace constituent is an internal component of the Infinite Volume.

What data constituents are

In an Infinite Volume, data is stored in multiple separate data constituents. Data constituents store only the data from a file, not the file's name.

Clients are not aware of data constituents. When a client requests a file from an Infinite Volume, the node retrieves the file's data from a data constituent and returns the file to the client.

Each Infinite Volume typically has dozens of data constituents. For example, a 6 PB Infinite Volume that contains 1 billion files might have 60 data constituents located on aggregates from 6 nodes.

What a namespace mirror constituent is

A namespace mirror constituent is an intracluster data protection mirror copy of the namespace constituent in an Infinite Volume. The namespace mirror constituent performs two roles: It provides data protection of the namespace constituent, and it supports SnapDiff for incremental tape backup of Infinite Volumes.

Creating rules

You can add new rules to your data policy to determine the placement of data that is written to the Infinite Volume. You can create rules either by using rule templates that are defined in Unified Manager or create custom rules.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- 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 299

• *Creating custom rules* on page 300

Related tasks

Adding a user on page 316

Creating rules using templates

You can add new rules by using rule templates defined by Unified Manager to determine the placement of data that is written to the Storage Virtual Machine (SVM) with Infinite Volume. You can create rules based on file types, directory paths, or owners.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click the **Data Policy** tab.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Click **Create**.
5. In the **Create Rule** dialog box, choose an appropriate rule template from the drop-down list.

   The template is based on three categories: file type, owner, or directory path.
6. Based on the template selected, add necessary conditions in the **Matching Criteria** area.
7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.
8. Click **Create**.

   The new rule you created is displayed in the Data Policy tab.
9. Optional: Preview any other changes made to the data policy.
10. Click **Activate** to activate the changes in the rule properties in the SVM.

Related concepts

What a rule template is on page 309
What conditions and condition sets are on page 309

Related tasks

Adding a user on page 316

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 be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage** > **Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.
4. Click **Create**.
5. In the **Create Rule** dialog box, select **Custom rule** from the **Template** list.
6. In the **Matching Criteria** area, add conditions as required.

   Conditions enable you to create a rule based on file types, directory paths, or owners. A combination of these conditions are the condition sets. For example, you can have a rule: “Place all .mp3 owned by John in bronze storage class.”
7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.
8. Click **Create**.

   The newly created rule is displayed in the Data Policy tab.
9. Optional: Preview any other changes made to the data policy.

10. Click **Activate** to activate the changes in the rule properties in the SVM.

Related concepts

What conditions and condition sets are on page 309

Related tasks

Adding a user on page 316

Viewing rules

You can view the list of rules you created from the Data Policy tab before modifying the data policy for your SVM with Infinite Volume.

Before you begin

- You must be assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.

Result

The list of rules in the data policy for the selected SVM is displayed. You can use Filter by Storage Class to view rules about a specific storage class.

Related concepts

What rules and data policies are on page 307

Related tasks

Adding a user on page 316
Creating rules on page 298

Editing template-based rules

You can edit a rule that was created using the rule templates from the Edit Rule dialog box. You can add, modify, or delete rule properties such as directory paths, file types, and owners. You can also modify the rule name and the storage class associated with the rule.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate Storage Virtual Machine (SVM).
3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Select the rule for which you want to include new conditions or condition sets.
5. Click **Edit**.
6. In the **Edit Rule** dialog box, edit the rule as required:

If you want to...	Do this...
Add a new rule property	Click **Add**.
Delete a rule property	Click **Delete** by selecting the appropriate rule property.
Modify a rule property	Double-click the appropriate rule property, and then modify as required.

7. Click **Update**.
8. Verify that your modifications are applied to the rule by expanding the rule in the **Data Policy** tab.
9. Optional: Preview any other changes made to the data policy.
10. Click **Activate** to activate the changes to the rule properties in the SVM.

Related concepts

What a rule template is on page 309

Related tasks

Adding a user on page 316

Editing custom rules

You can edit a rule to include new conditions or condition sets in the rule. For example, if you want to include new directory paths along with the owner names, you can do so from the Edit Rule dialog box.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

The Data Policy tab is visible only for a Storage Virtual Machine (SVM) with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Select the rule for which you want to include new conditions or condition sets.
5. Click **Edit**.
6. In the **Edit Rule** dialog box, add new conditions or condition sets:

If you want to add...	Click...
A new condition	The icon.
A new condition set	**Add Condition Set**.

7. Click **Update**.
8. Verify that your modifications are applied to the rule by expanding the rule in the **Data Policy** tab.
9. Optional: Preview any other changes made to the data policy.
10. Click **Activate** to activate the changes in the rule properties in the SVM.

Related concepts

What conditions and condition sets are on page 309

Related tasks

Adding a user on page 316

Deleting rules

You can delete a rule from a data policy when it is no longer required. For example, you might want to delete a rule on a particular directory that is no longer valid.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Select the rule that you want to delete, and then click **Delete**.

   Note: You cannot delete the default rule.
5. Optional: Preview any other changes made to the data policy.
6. Click **Activate** to activate the changes in the rule properties in the SVM.

Related tasks

Adding a user on page 316

Previewing changes to your data policy

You should preview any changes that you have made to your rules in a data policy before you submit the changes in the data policy to the Storage Virtual Machine (SVM) with Infinite Volume for activation.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Modify the data policy as required.

   Data policy modifications can include creating new rules, editing existing rules, deleting existing rules, or reordering the rules.
5. Click **Activate**.
6. In the **Summary of Changes to Data Policy Configuration** window, preview the changes to your data policy, and then click **Activate** to activate the changes in the data policy in the SVM with Infinite Volume.

Related tasks

Adding a user on page 316

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 be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

The Data Policy tab, which is used while performing this task, is displayed only for SVMs with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.

3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

4. Click **Export**.

5. In the browser-specific dialog box, specify the location to which the data policy configuration has to be exported.

Result

The data policy configuration is exported as a JSON file in the specified location.

Related tasks

Importing a data policy configuration on page 306
Adding a user on page 316

Importing a data policy configuration

You can import a data policy configuration from a file, modify the data policy, and then activate the changes in the SVM with Infinite Volume. For example, in the event of a disaster, you can import an already defined data policy to the secondary and modify the policy as required.

Before you begin

You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

About this task

When you import a data policy configuration, your existing rules are overwritten.

The Data Policy tab is displayed only for SVMs with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.

2. In the Storage Virtual Machines page, select the appropriate SVM.

3. Click **Data Policy**.

   The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

4. Click **Import**.

5. In the **Import Data Policy** dialog box, specify the data policy that you want to import by providing the absolute path of the data policy file.

6. Click **Import**.

7. Click **Activate** to activate the imported rules in the SVM.

Related tasks

Exporting a data policy configuration on page 305
Adding a user on page 316

Understanding rules and data policy

Understanding the concepts about rules and data policy help you to manage your Infinite Volumes efficiently.

What rules and data policies are

A *rule* determines the placement of files (data) in a Storage Virtual Machine (SVM) with Infinite Volume. A collection of such rules is known as a *data policy*.

Rule

Rules mainly consist of a set of predefined conditions and information that determine where to place files in the Infinite Volume. When a file is placed in the Infinite Volume, the attributes of that file are matched with the list of rules. If attributes match the rules, then that rule's placement information determines the storage class where the file is placed. A default rule in the data policy is used to determine the placement of files if the attributes do not match any of the rules in the rule list.

For example, if you have a rule, “Place all files of type .mp3 in the bronze storage class.”, all .mp3 files that are written to the Infinite Volume would be placed in the bronze storage class.

Data policy

A data policy is a list of rules. Each SVM with Infinite Volume has its own data policy. Each file that is added to the Infinite Volume is compared to its data policy's rules to determine where to place that file. The data policy enables you to filter incoming files based on the file attributes and place these files in the appropriate storage classes.

Related concepts

What the default rule is on page 308
How a data policy filters data written to an Infinite Volume on page 308

Related tasks

Creating rules on page 298

What the default rule is

The default rule is the rule present in the data policy of a Storage Virtual Machine (SVM) with Infinite Volume. It is used to determine the placement of data written to the Infinite Volume when none of the conditions in the existing rules match with the data being written.

The default rule is always the last rule in a data policy and cannot be reordered. For example, consider a data policy with three rules. Rule-1 places all .pdf files in the *high_performance* storage class. Rule-2 places all files owned by the administrator and file names that end with *.xls in the *archival_constituent* storage class. The third rule is the default rule with the *low_performance* storage class.

When a set of *.jpg files that are not owned by the administrator is written to the Infinite Volume, the default rule is used to place these .jpg files in the *low_performance* storage class. Rule-1 and Rule-2 are not used because the data that is written does not match these rules.

Related concepts

What rules and data policies are on page 307

How a data policy filters data written to an Infinite Volume

A data policy automatically filters data written to the Infinite Volume into different storage classes. All files are written to the single file system in the Infinite Volume's namespace, and rules in the data policy determine which storage class stores the data for the files.

A data policy with a default rule is automatically created for a Storage Virtual Machine (SVM) with Infinite Volume when you create the Infinite Volume. The data policy is active and contains a default rule. The default rule stores incoming data for files as follows for Infinite Volumes with and without storage classes:

For an Infinite Volume...	The default rule does this...
Without storage classes	Places all incoming data for files in the Infinite Volume
With one storage class	Places all incoming data for files into the storage class

For an Infinite Volume...	The default rule does this...
With one or more storage classes	Places all incoming data for files into the first storage class that is created **Important:** For an Infinite Volume with two or more storage classes, you should modify the data policy as soon as possible to create rules that filter data for different types of files into the different storage classes. You should modify the data policy by using OnCommand Unified Manager.

Note: The data policy does not affect the location of the files in the file system in the Infinite Volume's namespace, and client applications are unaware of storage classes. The file system in the namespace contains the file names. The data policy affects only which storage class is used to store the data for the files. Data policies are useful when you assign two or more storage classes to an Infinite Volume.

You can modify the data policy to create additional rules, but you cannot delete the data policy or its default rule.

What a rule template is

A rule template is a predefined template that can be used to create rules in a data policy. A rule template enables you to create a rule based on three categories: owner, file type, and directory path.

Example of a rule template for file types

The rule template “Place all files with the specified extensions in a suitable storage class” places all the .mp3 files that are written to the Infinite Volume in a storage class that you specify.

Related concepts

What rules and data policies are on page 307

What conditions and condition sets are

Conditions are a set of matching criteria using rule properties—such as the file name, directory path, and owner—that define a rule. A collection of such conditions is known as *condition set*. You can

use conditions and condition sets only for custom rules to determine where to place content that is written into your Infinite Volume.

Conditions

For a custom rule, you can specify conditions based on rule properties such as the file name, directory path, or owner, or a combination of all the rule properties. The logic is similar to a boolean AND operation. For example, by using conditions, you can create a custom rule to place files with .mp3 extensions and files owned by John in the directory path starting with `/NS/`.

Condition sets

The logic used for condition sets is similar to a boolean OR operation. For example, by using conditions and condition sets, you can create a complex custom rule that matches either of the following conditions:

- condition-1
  All files owned by Mary and are placed in `/NS/Engg/`
- condition-2
  All files that have names ending with .pdf and owned by Mary

Description of Infinite Volume windows and dialog boxes

You can monitor SVMs with Infinite Volume from the respective Storage Virtual Machine details page. You can manage rules and data policies from the Create Rule dialog box. You can also modify the storage class thresholds from the Edit Storage Class Thresholds dialog box.

Create Rule dialog box

You can use the Create Rule dialog box to create new rules for your data policy. For example, when you want to specify the placement of content of a certain file type, you can use the Create Rule dialog box to create the rule for your data policy.

- *Rule Name* on page 310
- *Templates area* on page 311
- *Matching Criteria* on page 311
- *Content Placement* on page 311
- *Command buttons* on page 311

Rule Name

Specifies the name of the new rule.

Templates area

Displays the list of rule templates. You can select an appropriate rule template from the list to create a rule for the data policy.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected. For example, if you select “Place all files with the specified owner names in a suitable storage class”, **List of Owner that...** is displayed in Matching Criteria.

Add Enables you to add a new rule property based on the rule template selected. For example, if you selected the rule template, “Place all files with the specified owner names in a suitable storage class”, the **Add** button enables you to add the owner's name.

Delete Enables you to delete a selected rule property.

Content Placement

Enables you to select an appropriate storage class for your rule from the list.

Command buttons

Create Creates a new rule for the data policy and closes the Create Rule dialog box.

Cancel Cancels the recent changes made to the rule and closes the Create Rule dialog box.

Related tasks

Creating rules on page 298

Edit Rule dialog box

You can use the Edit Rule dialog box to edit the properties of a rule, such as the file types, directory paths, or owners. You can also select an appropriate storage class for the rule. For example, when a certain file path is no longer valid, you can delete the file path from the corresponding rule.

- *Rule Name* on page 311
- *Matching Criteria* on page 312
- *Content Placement area* on page 312
- *Command buttons* on page 312

Rule Name

Displays the name of the rule.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected.

Add Enables you to add a new rule property, a new file type, a file path, or a new owner. For example, if you had specified the rule template, “Place all files with the specified owner names in a suitable storage class”, the Add button enables you to add the owner's name.

Delete Enables you to delete a selected rule property.

Content Placement area

Displays the list of storage classes. You can select an appropriate storage class for the selected rule.

Command buttons

Update Updates the changes made to the rule and closes the Edit Rule dialog box.

Cancel Cancels the recent changes made to the rule and closes the Edit Rule dialog box.

Related tasks

Editing template-based rules on page 302
Editing custom rules on page 303

Edit Rule dialog box (Advanced edit)

You can use the Edit Rule dialog box to edit the properties of a rule that is not created by using a template. The rule properties you can edit include the file types, directory paths, matching criteria, or owners. You can select an appropriate storage class for the rule. For example, you can edit the conditions specified in the matching criteria of a rule.

- *Rule Name* on page 312
- *Matching Criteria* on page 312
- *Content Placement area* on page 313
- *Command buttons* on page 313

Rule Name

Displays the name of the rule.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected. You can expand the rules and modify the rule properties, as required.

Content Placement area

Displays the list of storage classes. You can select an appropriate storage class for the selected rule.

Command buttons

Update	Updates the changes made to the rule and closes the Edit Rule dialog box.
Cancel	Cancels the recent changes made to the rule and closes the Edit Rule dialog box.

Edit SVM with Infinite Volume Thresholds dialog box

You can use the Edit SVM with Infinite Volume Thresholds dialog box to modify the default threshold values of each SVM with Infinite Volume based on your organization's requirements. The default threshold values indicate the level of activity that must be reached on the SVM before an event is triggered.

Capacity

The Capacity area enables you to set capacity threshold conditions for the selected SVM with Infinite Volume:

Space Nearly Full	Specifies the percentage at which the SVM with Infinite Volume is considered to be nearly full. It also displays the corresponding space (in GB, MB, or TB) in the Infinite Volume. For example, if you have an Infinite Volume of size 10 GB, and the Space Nearly Full threshold is 80%, then the following information is displayed: (8 GB of 10 GB). You can also use the slider to set the threshold value.
Space Full	Specifies the percentage at which the SVM with Infinite Volume is considered full. It also displays the corresponding space (in GB, MB, or TB) in the Infinite Volume. For example, if you have an Infinite Volume of size 10 GB, and the Space Full threshold is 90%, then the following information is displayed: (9 GB of 10 GB). You can also use the slider to set the threshold value.
Snapshot Usage Limit	Specifies the limit, in percentage, of space reserved for Snapshot copies in the Infinite Volume.

Command buttons

The command buttons enable you to perform the following tasks:

Restore to Global Defaults	Enables you to restore the threshold settings to the current values that are set at the global level.
Save	Saves all the threshold settings.

Save and Close	Saves all the threshold settings and then closes the Edit SVM with Infinite Volume Thresholds dialog box.
Cancel	Ignores any changes to the threshold settings and closes the Edit SVM with Infinite Volume Thresholds dialog box.

Related tasks

Editing the Infinite Volume threshold settings on page 295

Edit Storage Class Thresholds dialog box

You can use the Edit Storage Class Thresholds dialog box to modify the default threshold values of various storage classes in each SVM with Infinite Volume based on your organization's requirements. The default threshold values indicate the level of activity that must be reached on a storage class before an event is triggered.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Capacity

The Capacity area enables you to set capacity threshold conditions for the selected storage class.

Space Nearly Full	Specifies the percentage at which a storage class in the SVM with Infinite Volume is considered to be nearly full. It also displays the corresponding space (in GB, MB, or TB) in the storage class. For example, if you have a storage class of size 10 GB and the Space Nearly Full threshold is 80%, then the following information is displayed: (8 GB of 10 GB). You can also use the slider to set the threshold value.
Space Full	Specifies the percentage at which the storage class in the SVM with Infinite Volume is considered full. It also displays the corresponding space (in GB, MB, or TB) in the storage class. For example, if you have a storage class of size 10 GB and the Space Full threshold is 90%, then the following information is displayed: (9 GB of 10GB). You can also use the slider to set the threshold value.
Snapshot Usage Limit	Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.

Command buttons

The command buttons enable you to perform tasks for a selected volume.

Restore to Global Defaults	Enables you to restore the threshold settings to the current values that are set at the global level.
Save	Saves all the threshold settings.
Save and Close	Saves all the threshold settings and then closes the Edit Storage Class Thresholds dialog box.
Cancel	Cancels changes (if any) to the threshold settings and closes the Edit Storage Class Thresholds dialog box.

Related tasks

Editing the threshold settings of storage classes on page 296

Managing user access

You can create roles and assign capabilities to control user access to selected cluster objects. You can identify users who have the required capabilities to access selected objects within a cluster. Only these users are provided access to manage the cluster objects.

Adding a user

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

Before you begin

- To add a remote user or group, you must have enabled remote authentication and configured your authentication server.
- You must be logged in as the OnCommand Administrator role to perform this task.

About this task

If you add a group from active directory, then all direct members and nested subgroups can authenticate to Unified Manager. If you add a group from OpenLDAP or Other authentication services, then only direct members of that group can authenticate to Unified Manager.

Steps

1. Click **Administration** > **Manage Users**.
2. In the **Manage Users** page, click **Add**.
3. In the **Add User** dialog box, select the type of user that you want to create and enter the required information.

   When entering the required user information, you must specify an email address unique to that user. Specifying email addresses shared by multiple users must be avoided.
4. Click **Add**.

Related tasks

Enabling remote authentication on page 327
Setting up authentication services on page 328
Adding authentication servers on page 329

Related references

Unified Manager roles and capabilities on page 322
Definitions of user types on page 321
Definitions of user roles in Unified Manager on page 321

Editing the user settings

You can edit user settings, such as the email address and role that are specified for users on the Manage Users page. For example, you might want to change the role of a user who is a storage operator and assign storage administrator privileges to the user.

Before you begin

You must be logged in as the OnCommand Administrator role to perform this task.

About this task

When you modify the role assigned to a user, the changes are applied when either of the following occurs:

- The user logs out and logs back in to Unified Manager
- Session timeout of 24 hours

Steps

1. Click **Administration** > **Manage Users**.
2. In the **Manage Users** page, select the user that you want to edit and click **Edit**.
3. In the **Edit User** dialog box, edit the appropriate settings that are specified for the user.
4. Click **Save**.

Related tasks

Adding a user on page 316

Testing a remote user or group

You can validate a remote user or group by using the authentication settings specified for your authentication servers.

Before you begin

- You must have enabled remote authentication and configured your authentication settings so that the OnCommand Unified Manager server can validate the remote user or group.
- You must be logged in as the OnCommand Administrator role to perform this task.

Steps

1. Click **Administration > Manage Users**.
2. In the **Manage Users** page, select a remote user or group that you want to validate, and then click **Test**.

Related tasks

Adding a user on page 316
Enabling remote authentication on page 327
Setting up authentication services on page 328
Testing the configuration of authentication servers on page 330

Viewing users

You can use the Manage Users page to view the list of users who manage storage objects and data using Unified Manager. You can view details about the users, such as the name, type of user, email address, and role assigned to the users.

Before you begin

You must be logged in as the OnCommand Administrator role to perform this task.

Step

1. Click **Administration > Manage Users**.

   The list of users is displayed in the Manage Users page.

Related tasks

Adding a user on page 316

Deleting users or groups

You can delete one or more users from the management server database to prevent the users from accessing Unified Manager.

Before you begin

- When you are deleting remote groups, you must have reassigned the events that are assigned to users of the remote groups.

  However, if you are deleting local users or remote users, events assigned to these users are automatically unassigned.
- You must be logged in as the OnCommand Administrator role to perform this task.

Steps

1. Click **Administration** > **Manage Users**.
2. In the **Manage Users** page, select the users or groups that you want to delete and click **Delete**.
3. Click **Yes** to confirm the delete request.

Related tasks

Adding a user on page 316
Assigning events on page 52

Changing the local user password

You can change your login password to prevent potential security risks.

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. To change the maintenance user password, use the Unified Manager maintenance console. To change the remote user password, contact your password administrator.

Steps

1. Log in to Unified Manager.
2. Click ***user_name*** > **Change Password**.

   The **Change Password** option is not displayed if you are a remote user.
3. In the **Change Password** dialog box, enter the details as required.
4. Click **Save**.

Related tasks

Adding a user on page 316

What the maintenance user does

Created during initial configuration, the maintenance user can create subsequent users and assign them roles. The maintenance user can also access the maintenance console and has the role of OnCommand administrator in the web UI.

The maintenance user can perform the following functions using the maintenance console:

- Configure network access
- Upgrade to newer versions of Unified Manager
- Shut down virtual appliances (only from VMware console)
- Increase data disk or swap disk size
- Change the time zone
- Send on-demand AutoSupport messages to technical support from the maintenance console
- Send periodic AutoSupport messages to technical support from the web UI
- Generate support bundles to send to technical support

Related references

Definitions of user types on page 321
Definitions of user roles in Unified Manager on page 321
Unified Manager roles and capabilities on page 322

What RBAC is

RBAC (role-based access control) provides the ability to control who has access to various features and resources in the OnCommand Unified Manager server.

What role-based access control does

Role-based access control (RBAC) enables administrators to manage groups of users by defining roles. If you need to restrict access for specific functionality to selected administrators, you must set up administrator accounts for them. If you want to restrict the information that administrators can view and the operations they can perform, you must apply roles to the administrator accounts you create.

The management server uses RBAC for user login and role permissions. If you have not changed the management server's default settings for administrative user access, you do not need to log in to view them.

When you initiate an operation that requires specific privileges, the management server prompts you to log in. For example, to create administrator accounts, you must log in with Administrator account access.

Definitions of user types

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

Unified Manager user types are as follows:

Maintenance user	Created from the maintenance console 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.
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	Groups of users that access the Unified Manager UI using the credentials stored on the authentication server. The name of this account should match the name of a group stored on the authentication server. All users within the remote group are given access to the Unified Manager UI using their individual user credentials. Remote groups can perform functions according to their assigned roles.
Remote user	Accesses the Unified Manager UI using the credentials stored on the authentication server. A remote user performs functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.
Database user	Has read-only access to data in the Unified Manager database, has no access to the Unified Manager web interface or the maintenance console, and cannot execute API calls.

Related concepts

What the maintenance user does on page 320

Definitions of user roles in Unified Manager

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

The following predefined roles exist in Unified Manager:

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, 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.
Event Publisher	Transmits events generated by partner applications for display in the Unified Manager web UI. This specialized, limited-access user role enables partner applications to share event information with Unified Manager. At the same time, the limited access of this role prevents unauthorized access to the Unified Manager server or the Unified Manager web UI through event publishing activity.

Related concepts

What the maintenance user does on page 320

Unified Manager roles and capabilities

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

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

Function	Event Publisher	Operator	Storage Administrator	OnCommand Administrator
View storage system information		•	•	•
View other data like histories, capacity trends, and so on		•	•	•
View, assign, resolve events		•	•	•
View storage service objects, such as SVM associations and resource pools		•	•	•

Function	Event Publisher	Operator	Storage Administrator	OnCommand Administrator
Manage storage service objects, such as SVM associations and resource pools			•	•
Define alerts			•	•
Manage storage management options			•	•
Manage storage management policies			•	•
Manage users				•
Manage administrative options				•
Manage database access				•
Publish events	•			

Related concepts

What the maintenance user does on page 320

Description of user access windows and dialog boxes

Based on the RBAC settings, you can add users from the Manage Users page and assign appropriate roles to those users to access and monitor your clusters.

Manage Users page

The Manage Users page displays a list of your users and groups, and provides information such as the name, type of user, and email address. You can also perform tasks such as adding, editing, deleting, and testing users.

- *Command buttons* on page 323
- *List view* on page 324

Command buttons

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

Add Displays the Add User dialog box, which enables you to add a local user, remote user, remote group, or a database user.

You can add remote users or groups only if your authentication server is enabled and configured.

Edit Displays the Edit User dialog box, which enables you to edit the settings for the selected user.

Delete Deletes the selected users from the management server database.

Test Enables you to validate whether a remote user or group is present in the authentication server.

You can perform this task only if your authentication server is enabled and configured.

List view

The List view displays, in tabular format, information about the users that are created. You can use the column filters to customize the data that is displayed.

Name Displays the name of the user or group.

Type Displays the type of user. The user type can be Local User, Remote User, Remote Group, Database User, or Maintenance User.

Email Displays the email address of the user.

Role Displays the type of role that is assigned to the user. The role can be Operator, Storage Administrator, or OnCommand Administrator.

Related tasks

Viewing users on page 318
Adding a user on page 316
Editing the user settings on page 317
Testing a remote user or group on page 317
Deleting users or groups on page 318

Add User dialog box

You can create local users or database users, or add remote users or remote groups and assign roles so that these users can efficiently manage the storage objects and data using Unified Manager.

You can add a user by completing the following fields:

Type Enables you to specify the type of user you want to create.

Name Enables you to specify a user name that a user can use to log in to Unified Manager.

Password	Enables you to specify a password for the specified user name. This field is displayed only when you are adding a local user or a database user.
Confirm Password	Enables you to reenter your password to ensure the accuracy of what you entered in the Password field. This field is displayed only when you are adding a local user or a database user.
Email	Enables you to specify an email address for the user. This field is displayed only when you are adding a remote user or a local user. **Note:** The email address specified must be unique to the user name.
Role	Enables you to assign a role to the user and defines the scope of activities that the user can perform. The role can be OnCommand Administrator, Storage Administrator, or Operator.

Command buttons

The command buttons enable you to perform the following tasks:

Add	Adds the user and closes the Add User dialog box.
Cancel	Cancels the changes and closes the Add User dialog box.

Related tasks

Adding a user on page 316

Edit User dialog box

The Edit User dialog box enables you to edit the email address or role of a selected user.

Details

The Details area enables you to edit the following information about a selected user:

Type	Enables you to modify the type of user.
Name	Enables you to change the user name of the selected user.
Email	Enables you to edit the email address of the selected user.
Role	Enables you to edit the role that is assigned to the user. This field is displayed only when the selected user is a local user, remote user, or remote group.

Command buttons

The command buttons enable you to perform the following tasks:

Save	Saves the changes and closes the Edit User dialog box.

Cancel Cancels the changes and closes the Edit User dialog box.

Related tasks

Editing the user settings on page 317

Managing authentication

You can enable LDAP authentication on the Unified Manager server and configure it to work with your LDAP servers to authenticate remote users.

Enabling remote authentication

Using Open LDAP or Active Directory, you can enable remote authentication so that the management server can communicate with your authentication servers and so that users of the authentication servers can use Unified Manager to manage the storage objects and data.

Before you begin

You must be logged in as the OnCommand Administrator role to perform this task.

About this task

If remote authentication is disabled, remote users or groups can no longer access Unified Manager.

The only two supported remote authentication methods are Active Directory and Open LDAP. LDAPS is not supported.

Steps

1. Click **Administration** > **Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server** > **Authentication**.
3. Select **Enable Remote Authentication**.

   If you are using Active Directory Authentication Service, you can enter the Administrator Name using one of the following formats:

   - domainname/username
   - username@domainname
   - Bind Distinguished Name (using appropriate LDAP notation)
4. Optional: Add authentication servers and test the authentication.
5. Click **Save and Close**.

Related tasks

Adding a user on page 316

Setting up authentication services

Authentication services enable the authentication of remote users or groups in an authentication server before providing them access to Unified Manager. You can authenticate users by using the predefined authentication services, such as Active Directory or OpenLDAP, or by configuring your own authentication mechanism.

Before you begin

- You must have enabled remote authentication.
- You must be logged in as the OnCommand Administrator role to perform this task.

Steps

1. Click **Administration** > **Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server** > **Authentication**.
3. Select one of the following authentication services:

If you select...	Then do this...
Active Directory	**a.** Enter the administrator name and password. **b.** Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is `dc=ou,dc=domain,dc=com`.
OpenLDAP	**a.** Enter the bind distinguished name and bind password. **b.** Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is `dc=ou,dc=domain,dc=com`.
Others	**a.** Enter the bind distinguished name and bind password. **b.** Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is `dc=ou,dc=domain,dc=com`. **c.** Specify the LDAP protocol version that is supported by the authentication server. **d.** Enter the user name, group membership, user group, and member attributes.

Note: If you want to modify the authentication service, ensure that you delete any existing authentication servers and add new authentication servers.

4. Click **Save and Close**.

Related tasks

Adding a user on page 316
Enabling remote authentication on page 327

Adding authentication servers

You can add authentication servers and enable remote authentication on the management server to enable remote users within the authentication server to access Unified Manager.

Before you begin

- The following information must be available:
  - Host name or IP address of the authentication server
  - Port number of the authentication server
- You must have enabled remote authentication and configured your authentication service so that the management server can authenticate remote users or groups in the authentication server.
- You must be logged in as the OnCommand Administrator role to perform this task.

About this task

If the authentication server that you are adding is part of a high-availability pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > Authentication**.
3. In the Servers area, click **Add**.
4. In the **Add Authentication Server** dialog box, specify either the host name or IP address of the server, and the port details.
5. Click **Add**.

Result

The authentication server that you added is displayed in the Servers area.

After you finish

Perform a test authentication to confirm that you are able to authenticate users in the authentication server that you added.

Related concepts

Authentication with Active Directory or OpenLDAP on page 332

Related tasks

Adding a user on page 316
Enabling remote authentication on page 327
Setting up authentication services on page 328
Testing the configuration of authentication servers on page 330

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 the authentication servers. You can test by searching for a remote user or group from your authentication servers and authenticate them using the configured settings.

Before you begin

- You must have enabled remote authentication and configured your authentication service so that the OnCommand Unified Manager server can authenticate the remote user or group.
- You must have added your authentication servers so that the management server can search for the remote user or group from these servers and authenticate them.
- You must be assigned the OnCommand Administrator role to perform this task.

About this task

If the authentication service is set to Active Directory and if you are testing the authentication of remote users who belong to the primary group of the authentication server, information about the primary group is not displayed in the authentication results.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > Authentication**.
3. In the **Authentication Setup Options** dialog box, click **Test Authentication**.
4. In the **Test User** dialog box, specify the user name and password of the remote user or group, and then click **Test**.

   If you are authenticating a remote group, you must not enter the password.

Related tasks

Adding a user on page 316
Enabling remote authentication on page 327

Setting up authentication services on page 328
Adding authentication servers on page 329

Editing authentication servers

You can change the port that the OnCommand Unified Manager server uses to communicate with your authentication server.

Before you begin

You must be logged in as the OnCommand Administrator role to perform this task.

Steps

1. Click **Administration** > **Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server** > **Authentication**.
3. In the **Servers** area, select the authentication server that you want to edit, and then click **Edit**.
4. In the **Edit Authentication Server** dialog box, edit the port details.
5. Click **Save**.

Related tasks

Adding a user on page 316

Deleting authentication servers

You can delete an authentication server if you want to prevent the Unified Manager server from communicating with the authentication server. For example, if you want to change an authentication server that the management server is communicating with, you can delete the authentication server and add a new authentication server.

Before you begin

You must be logged in as the OnCommand Administrator role to perform this task.

About this task

When you delete an authentication server, remote users or groups of the authentication server will no longer be able to access Unified Manager.

Steps

1. Click **Administration** > **Setup Options**.

2. In the **Setup Options** dialog box, click **Management Server > Authentication**.
3. In the **Servers** area, select one or more authentication servers that you want to delete, and then click **Delete**.
4. Click **Yes** to confirm the delete request.

Related tasks

Adding a user on page 316

Authentication with Active Directory or OpenLDAP

You can enable remote authentication on the management server and configure the management server to communicate with your authentication servers so that users within the authentication servers can access Unified Manager.

You can use one of the following predefined authentication services or specify your own authentication service to:

- Microsoft Active Directory
- OpenLDAP

You can select the required authentication service and add the appropriate authentication servers to enable the remote users in the authentication server to access Unified Manager. The credentials for remote users or groups are maintained by the authentication server. The management server uses the Lightweight Directory Access Protocol (LDAP) to authenticate remote users within the configured authentication server.

For local users who are created in Unified Manager, the management server maintains its own database of user names and passwords. The management server performs the authentication and does not use Active Directory or OpenLDAP for authentication.

Description of authentication windows and dialog boxes

You can enable LDAP authentication from the Authentication Setup Options dialog box.

Authentication Setup Options dialog box

You can use the Authentication Setup Options dialog box to configure the management server to communicate with your authentication server and authenticate remote users in the authentication server.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

Enable Remote Authentication area

The Enable Remote Authentication area allows you to enable or disable remote authentication. You can enable remote authentication to enable the management server to authenticate remote users within the configured authentication servers.

Authentication Service

Enables you to configure the management server to authenticate users in directory service providers, such as Active Directory, OpenLDAP, or specify your own authentication mechanism. You can specify an authentication service only if you have enabled remote authentication.

Active Directory

- Administrator Name
  Specifies the administrator name of the authentication server.
- Password
  Specifies the password to access the authentication server.
- Base Distinguished Name
  Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is `dc=ou,dc=domain,dc=com`.

OpenLDAP

- Bind Distinguished Name
  Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server.
- Bind Password
  Specifies the password to access the authentication server.
- Base Distinguished Name
  Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is `dc=ou,dc=domain,dc=com`.

Others

- Bind Distinguished Name
  Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server that you have configured.
- Bind Password
  Specifies the password to access the authentication server.
- Base Distinguished Name
  Specifies the location of the remote users in the authentication server. For example, if the domain name of

the authentication server is ou@domain.com, then the base distinguished name is `dc=ou,dc=domain,dc=com`.

- Protocol Version

  Specifies the Lightweight Directory Access Protocol (LDAP) version that is supported by your authentication server. You can specify whether the protocol version must be automatically detected or set the version to 2 or 3.
- User Name Attribute

  Specifies the name of the attribute in the authentication server that contains user login names to be authenticated by the management server.
- Group Membership Attribute

  Specifies a value that assigns the management server group membership to remote users based on an attribute and value specified in the user's authentication server.
- UGID

  If the remote users are included as members of a GroupOfUniqueNames object in the authentication server, this option enables you to assign the management server group membership to the remote users based on a specified attribute in that GroupOfUniqueNames object.
- Member

  Specifies the attribute name that your authentication server uses to store information about the individual members of a group.
- Group Object Class

  Specifies the object class of all groups in the remote authentication server.

Note: If you want to modify the authentication service, ensure that you delete any existing authentication servers and add new authentication servers.

Servers area

The Servers area displays the authentication servers that the management server communicates with to find and authenticate remote users. The credentials for remote users or groups are maintained by the authentication server.

Command buttons

Enables you to add, edit, or delete authentication servers.

- Add

  Enables you to add an authentication server.

If the authentication server that you are adding is part of an high-availability pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.

- Edit

  Enables you to edit the settings for a selected authentication server.
- Delete

  Deletes the selected authentication servers.

Name or IP Address	Displays the host name or IP address of the authentication server that is used to authenticate the user on the management server.
Port	Displays the port number of the authentication server.

Test Authentication area

The Test Authentication area enables you to test your configuration.

Test Validates the configuration of your authentication server by authenticating a remote user or group.

While testing, if you specify only the user name, the management server searches for the remote user in the authentication server, but does not authenticate the user. If you specify both the user name and password, the management server searches and authenticates the remote user.

You cannot test the authentication if remote authentication is disabled.

Command buttons

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

Restore to Factory Defaults	Enables you to restore the configuration settings to the factory default values.
Save	Saves the configuration settings for the selected option.
Save and Close	Saves the configuration settings for the selected option and closes the Setup Options dialog box.
Cancel	Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Enabling remote authentication on page 327
Setting up authentication services on page 328
Adding authentication servers on page 329

Editing authentication servers on page 331
Testing the configuration of authentication servers on page 330
Deleting authentication servers on page 331

Managing security certificates

You can configure HTTPS in the Unified Manager server to monitor and manage your clusters over a secure connection.

Viewing the HTTPS security certificate

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

Before you begin

You must be logged in as the OnCommand Administrator role to perform this task.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > HTTPS**.
3. Click **View HTTPS Certificate**.

   The Subject DN field should display the same host name or fully qualified domain name (FQDN) that is displayed in the Configure Network Settings dialog box. The IP addresses should also be the same in the certificate and in the network settings.

   To view more detailed information about the security certificate, you can view the connection certificate in your browser.

Related tasks

Adding a user on page 316
Generating an HTTPS security certificate on page 338
Downloading an HTTPS certificate signing request on page 340
Installing an HTTPS security certificate on page 341
Restarting the Unified Manager virtual machine on page 338

Restarting the Unified Manager virtual machine

You can restart the virtual machine from the maintenance console. You might need to restart after generating a new security certificate or if there is a problem with the virtual machine.

Before you begin

The virtual appliance must be powered on.

You must be logged in to the NetApp maintenance console as the maintenance user.

About this task

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

Steps

1. Access the maintenance console.
2. Select **System Configuration** > **Reboot Virtual Machine.**
3. Start the Unified Manager GUI from your browser and log in.

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 be logged in as the OnCommand Administrator role to perform this task.

About this task

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

Steps

1. Click **Administration** > **Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server** > **HTTPS**.

3. Click **Regenerate HTTPS Certificate**.

   Important: You must restart the Unified Manager virtual machine before the new certificate takes effect. You can use the **System Configuration** option in the NetApp maintenance console.

After you finish

After generating a new certificate, you can verify the new certificate information by viewing the HTTPS certificate.

If you need to reactivate performance monitoring updates from Performance Manager servers to the Unified Manager server, you must delete the connections that were invalidated by this task and reconfigure new connections.

Related tasks

Adding a user on page 316
Viewing the HTTPS security certificate on page 337
Downloading an HTTPS certificate signing request on page 340
Installing an HTTPS security certificate on page 341
Accessing the maintenance console using Secure Shell on page 339
Accessing the maintenance console using the vSphere VM console on page 340

Related references

Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring on page 348

Accessing the maintenance console using Secure Shell

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

Before you begin

You must have installed and configured Unified Manager.

You must be logged in as the Maintenance User to perform this task.

Steps

1. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance.
2. Log in to the maintenance console using your maintenance user name and password.

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

Accessing the maintenance console using the vSphere VM console

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

Before you begin

You must be the maintenance user. The virtual appliance must be powered on to access the maintenance console.

Steps

1. In vSphere Client, locate the Unified Manager virtual appliance.
2. Click the **Console** tab.
3. Click inside the console window to log in.
4. Log in to the maintenance console using your user name and password.

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

Downloading an HTTPS certificate signing request

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

Before you begin

You must have the OnCommand Administrator role to perform this task.

Steps

1. Click **Administration** > **Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server** > **HTTPS**.
3. Click **Download HTTPS Certificate Signing Request**.
4. Save the `<hostname>.csr` file.

After you finish

You can provide the file to a Certificate Authority to sign and then install the signed certificate.

Related tasks

Adding a user on page 316
Viewing the HTTPS security certificate on page 337
Generating an HTTPS security certificate on page 338
Installing an HTTPS security certificate on page 341

Installing an HTTPS security certificate

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

Before you begin

You must have completed the following actions:

- Downloaded the Certificate Signing Request file and had it signed by a Certificate Authority
- Saved the certificate chain in PEM format
- Included all certificates in the chain, from the server certificate to the root signing certificate

You must have the OnCommand Administrator role to perform this task.

About this task

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

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > HTTPS**.
3. Click **Install HTTPS Certificate**.

Example certificate chain

The following example shows how the certificate chain file might appear:

```
-----BEGIN CERTIFICATE-----
<Server certificate>
-----END CERTIFICATE-----
```

```
-----BEGIN CERTIFICATE-----
<Root signing certificate>
-----END CERTIFICATE-----
```

After you finish

If you need to reactivate performance monitoring updates from Performance Manager servers to the Unified Manager server, you must delete the connections that were invalidated by this task and reconfigure new connections.

Related tasks

Adding a user on page 316
Viewing the HTTPS security certificate on page 337
Generating an HTTPS security certificate on page 338
Downloading an HTTPS certificate signing request on page 340

Related references

Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring on page 348

Description of security certificates windows and dialog boxes

You can use the Setup Options dialog box to view the current security certificates and to generate new HTTPS certificates.

HTTPS Setup Options dialog box

The HTTPS area in the Setup Options dialog box enables you to view the current security certificate, download a certificate signing request, to generate a new HTTPS certificate, or install a new HTTPS certificate.

You can edit the settings only if you are assigned either the OnCommand Administrator role or the Storage Administrator role.

HTTPS Certificate

You can perform the following operations:

View HTTPS Certificate	Enables you to view the current HTTPS certificate. If you have not generated a new HTTPS certificate, this is the certificate that was generated with your installation.

Regenerate HTTPS Certificate	Enables you to generate an HTTPS certificate, which replaces the previous security certificate. The new certificate is in effect after you restart the management server.
Download HTTPS Certificate Signing Request	Downloads a certification request for the currently installed HTTPS certificate. Your browser prompts you to save the `<hostname>.csr` file so that you can provide the file to a Certificate Authority to sign.
Install HTTPS Certificate	Enables you to upload and install a security certificate after a Certificate Authority has signed and returned it. The new certificate is in effect after you restart the management server.

Command buttons

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

Restore to Factory Defaults	Enables you to restore the configuration settings to the factory default values.
Save	Saves the configuration settings for the selected option.
Save and Close	Saves the configuration settings for the selected option and closes the Setup Options dialog box.
Cancel	Cancels the recent changes and closes the Setup Options dialog box.

Related tasks

Viewing the HTTPS security certificate on page 337
Generating an HTTPS security certificate on page 338
Downloading an HTTPS certificate signing request on page 340
Installing an HTTPS security certificate on page 341
Working with HTTPS security certificates on page 27

Troubleshooting

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

Sending an on-demand AutoSupport message

You can choose to have OnCommand Unified Manager send an on-demand message to technical support for assistance with troubleshooting issues. The AutoSupport message sent by Unified Manager contains diagnostic system information and detailed data about the Unified Manager server.

Before you begin

You must be logged in as the Maintenance User to perform this task.

Steps

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

   The dialog box also displays the product serial number, which is the number that technical support uses to find the AutoSupport messages sent from Unified Manager.
4. Click **Generate and Send AutoSupport**.

Related tasks

Adding a user on page 316
Enabling periodic AutoSupport on page 25

Unknown authentication error

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

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

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

Corrective action

1. Click **Administration** > **Setup Options**.
2. In the Setup Options dialog box, click **Management Server** > **Authentication**.
3. Based on the authentication service that you have selected, enter the appropriate information for Administrator Name or Bind Distinguished Name.
4. Click **Test** to test the authentication with the details you specified.
5. Click **Save and Close**.

User not found

Issue

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

Cause

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

Corrective action

1. Click **Administration** > **Setup Options**.
2. In the Setup Options dialog box, click **Management Server** > **Authentication**.
3. Enter the appropriate information for Base Distinguished Name.
4. Click **Save and Close**.

Icons are misaligned in Internet Explorer

Issue

Icons and text are misaligned when you use Internet Explorer.

Cause

This problem can occur if you are using Internet Explorer in Compatibility View, which is not a supported browser setting.

Corrective action

1. Press F12 to open Internet Explorer Developer Tools.
2. Select **Browser Mode** from the toolbar to display the browser version used to open the application.
3. Select **Document Mode** from the toolbar and select the Standards mode of the browser version used to open the application.

   For example, if you are using Internet Explorer 9 to open the application, select **Browser Mode** > **Internet Explorer 9**, and then select **Document Mode** > **Internet Explorer 9 Standards**.

LDAP server slow to respond

Issue The LDAP server takes a long time to respond to queries.

Cause Supporting nested groups causes the LDAP server to slow down.

Corrective action If you use Active Directory, you can speed authentication by disabling support for nested groups in Unified Manager. However, if you choose to disable nested groups, you must ensure that users are direct members of the groups added to Unified Manager.

To disable nested group support, follow these steps:

1. Click **Administration** > **Setup Options**.
2. In the Setup Options dialog box, click **Management Server** > **Authentication**.
3. In the **Authentication Service** drop-down menu, select **Others**.
4. In the **Member** box, type “member”.
5. Click **Save and Close**.

Issue with adding LDAP using Other authentication services

Issue When you select Other in the Authentication section of the Setup Options dialog box, the user and groupObjectClass retain the values from the previously selected template. If the LDAP server does not use the same values, the operation might fail.

Cause The users are not configured correctly in OpenLDAP.

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

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

1. Click **Administration** > **Setup Options**.
2. In the Setup Options dialog box, click **Management Server** > **Authentication**.
3. In the **Authentication Service** drop-down menu, select **Active Directory** and then select **Others**.
4. Complete the text fields.

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

1. Click **Administration** > **Setup Options**.

2. In the Setup Options dialog box, click **Management Server > Authentication**.
3. In the **Authentication Service** drop-down menu, select **OpenLDAP** and then select **Others**.
4. Complete the text fields.

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

Troubleshooting access to CIFS shares

You might not be able to access CIFS shares if the storage objects serving these shares are unavailable. You should review availability events such as Volume Offline, Junction Path Offline, or SVM CIFS Server Down that are generated when these objects are unavailable.

Before you begin

You must have the role of Storage Administrator to perform this task.

About this task

If you have configured an appropriate alert, you will be notified about the availability event through an alert email.

Steps

1. Log in to Unified Manager UI.
2. From the **Dashboard**, click the appropriate offline event.

   #### Example

   For example, if you receive a Volume Offline event, click the ***`Volume_name`*** **Volume Offline** event in the Availability panel in the Unresolved Incidents and Risks area.

3. In the **Event details** page, click ***`Volume_name`*** in the **Source** field.
4. In the **Volume details** page, click the number corresponding to CIFS Shares in the **Related Devices** pane.
5. In the **Storage Virtual Machine details** page, click the **CIFS Shares** tab.

   You can view the number of CIFS shares that are affected.

After you finish

You must resolve the failures by using either OnCommand System Manager or the Data ONTAP CLI.

Identifying inaccessible volumes and qtrees

You might not be able to access a volume or a qtree or write data to the volume or qtree. In such cases, the storage administrator must identify the quota for the volume or qtree that is not accessible and take appropriate action to resolve the issue.

Before you begin

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

Steps

1. Log in to Unified Manager.
2. In the search criteria, select **User or Group Quotas**, enter the user name, and press Enter.

   #### Example

   If a user with user name user1 cannot access a volume, select **User or Groups Quotas** in the search criteria, and then type `user1` or type the any three characters of the user name in the search bar.
3. Select the appropriate user from the drop-down list.

   The User and Group Quotas tab in the Storage Virtual Machine details page is displayed.

After you finish

You must resolve the issue by using either OnCommand System Manager or the Data ONTAP CLI.

Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring

Issue If performance monitoring is enabled on the Unified Manager server through connections with one or more Performance Manager servers, and if an HTTPS certificate is installed or regenerated on that Unified Manager server and that server is rebooted, posting of additional performance incidents encountered by Performance Manager servers to the Unified Manager Web UI is stopped.

Cause

The HTTPS certificate installation or regeneration on the Unified Manager server has invalidated the credentials that allowed the Performance Manager servers to post performance incidents to the Unified Manager server.

Corrective actions

You need to delete and then reconfigure the connections between the Unified Manager server and each Performance Manager server that was posting performance incidents to it before the HTTPS certificate installation or regeneration.

Before starting corrective actions, be prepared to respecify the current connection information:

- Unified Manager server name or IP address
- Unified Manager server port (always 443)
- Event Publisher user name (the name of the local Unified Manager server user assigned Event Publisher role privileges)
- Event Publisher password (the password of the local Unified Manager server user assigned Event Publisher role privileges)

Log in to the maintenance console of each Performance Manager server that was posting performance incidents to the Unified Manager server, and then complete the following actions:

1. Type the number of the menu option labeled "Unified Manager Connection" to display the Unified Manager Connection Menu, and then type the number of the menu option labeled "Delete Unified Manager Server Connection."
2. When prompted to confirm that you want to continue the deletion, type **y**, and then press any key to continue.
3. Type the number of the menu option labeled "Add / Modify Unified Manager Server Connection."
4. When prompted, supply the requested Unified Manager server name or IP address and Unified Manager server port information.
5. When prompted, accept the Unified Manager server trust certificate to support the connection. Do not change the default port value (443).
6. When prompted, supply the requested Event Publisher user name and Event Publisher password, and then confirm that the settings are correct.
7. To exit the maintenance console, press any key to continue, and then type **x**.

Glossary

A

Access Control List (ACL) A set of data associated with a file, directory, or other resource or share that defines user or group access rights to that resource or share.

admin SVM Formerly known as admin Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that has overall administrative access to all objects in the cluster, including all objects owned by other SVMs, but does not provide data access to clients or hosts.

aggregate A set of multiple RAID (Redundant Array of Inexpensive Disks or Redundant Array of Independent Disks) groups that can be managed as a single unit for protection and provisioning purposes.

aggregate committed capacity The data storage space in an aggregate that is committed to provide for its underlying volumes. Calculated by the total capacity provisioned for volumes.

aggregate total capacity The data storage space within an aggregate that can be used by volumes or aggregate-level Snapshot copies. Calculated by the total data capacity of the aggregate plus the aggregate-level Snapshot reserve space.

alert

- In OnCommand Insight (formerly SANscreen suite), an alarm indicating that a device or path state has changed in a way that violates a policy or exceeds a threshold.
- In Unified Manager, a user-configured notification that is sent whenever a specific event or an event of a specific severity type occurs, not necessarily related to a specific user. Alerts are used to monitor and manage datasets and resources. See also *event* and *severity type*.

AutoSupport An integrated technology that triggers email messages from the customer site to technical support or another specified email recipient when there are any failures in Unified Manager services. These messages contain information such as feature usage metrics, configuration and user settings, system health, and so on.

available capacity The amount of usable space available in a storage system. Calculated by the used capacity minus the unused reserve capacity.

B

backup relationship

A persistent association between a primary directory and a secondary volume for disk-based data backup and restore using the Data ONTAP SnapVault feature.

baseline transfer

An entire transfer of data as compared to an incremental transfer of data.

C

CIFS

See *Common Internet File System (CIFS)*.

CIFS share

- In Data ONTAP, a directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known simply as a *share*.
- In OnCommand Insight (formerly SANscreen suite), a service exposed from a NAS device to provide file-based storage through the CIFS protocol. CIFS is mostly used for Microsoft Windows clients, but many other operating systems can access CIFS shares as well.

client application

An application that calls Unified Manager APIs to enable its operator to configure, monitor, and initiate data management operations to be executed on the Unified Manager server.

cluster

- In clustered Data ONTAP 8.x, a group of connected nodes (storage systems) that share a namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.
- In the Data ONTAP 7.1 release family and earlier releases, a pair of storage systems (sometimes called *nodes*) configured to serve data for each other if one of the two systems stops functioning.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.
- For some storage array vendors, *cluster* refers to the hardware component on which host adapters and ports are located. Some storage array vendors refer to this component as a *controller*.

cluster committed capacity

The data storage space in a cluster that is committed to provide for its underlying aggregates. Calculated by the sum of the total capacity of all the aggregates in the cluster.

cluster failover (CFO)

In Data ONTAP 7.1.x and earlier, the method of ensuring data availability by transferring the data service of a failed node to another node in an HA pair.

Transfer of data service is often transparent to users and applications. In Data ONTAP 7.2 and later, and in Data ONTAP operating in 7-Mode, the failover method is called *controller failover*. In clustered Data ONTAP, the failover method is called *storage failover*.

cluster interconnect The cables and adapters with which two nodes (storage systems) in an HA pair are connected, and over which heartbeat and WAFL log information are transmitted when both nodes are running.

cluster total capacity The data storage space in a cluster that can be used by aggregates or volumes. Calculated by the sum of the capacity of all the data disks excluding disk right-sizing and reservation plus sum of the capacity of all spare disks excluding right-sizing.

cluster Vserver Former name for a data SVM; see data SVM.

container object An object, such as an aggregate or a Storage Virtual Machine (SVM, formerly known as Vserver), in which data objects reside.

counter The statistical measurement of activity on a storage system or storage subsystem that is provided by Data ONTAP. Each type of storage system or subsystem has a set of counters.

D

data capacity The storage space that is set aside by the container, such as an aggregate or a volume, to store user data. Typically, this capacity can be used for any container, but data is only written to the lowest level container, usually the volume.

Data ONTAP The operating system software running on NetApp storage devices.

datastore A storage location for virtual machines, such as a VMFS volume, a directory on a NAS server, or a local file system path. A datastore is platform-independent and host-independent; therefore, it does not change when a virtual machine it contains moves to another host.

data object A container of data, such as a file, directory, volume, or LUN, that can be discovered, monitored, protected, created, or restored by the Unified Manager server.

data SVM Formerly known as data Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data SVMs within a cluster.

deduplication The consolidation of blocks of duplicate data into single blocks to store more information using less storage space.

deduplication return The capacity savings resulting from deduplication. Calculated by the volume capacity before deduplication - the volume capacity after deduplication.

destination The storage to which source data is backed up, mirrored, or migrated.

destination data object A data object that contains the backed up or mirrored replicated data.

dedupe See *deduplication*.

DHCP See *Dynamic Host Configuration Protocol (DHCP)*.

Dynamic Host Configuration Protocol (DHCP) The protocol for automating the assignment of network addresses.

E

ESX server A VMware term describing a server that abstracts server processor, memory, storage, and networking resources into multiple virtual machines.

event An indication of a predefined condition occurring or when an object crosses a threshold. All events are assigned a severity type and are automatically logged in the Events window. See also *alert* and *severity type*.

F

failover The process by which an alternate storage system takes over and emulates a primary system if the primary system becomes unusable.

Fibre Channel (FC) A high-speed data transmission protocol, which is a licensed service on the storage system that enables you to export LUNs to hosts using the SCSI protocol over an FC fabric.

FlexVol volume A logical entity contained in a Storage Virtual Machine (SVM, formerly known as Vserver)—referred to as SVM with FlexVol volumes. FlexVol volumes typically hold user data, although they also serve as node or SVM root volumes and metadata containers. A FlexVol volume obtains its storage from a single aggregate.

FQDN See *Fully Qualified Domain Name (FQDN)*.

Fully Qualified Domain Name (FQDN) The complete name of a specific computer on the Internet, consisting of the computer's host name and its domain name.

fractional reserve An option that determines how much space in a volume is reserved for Snapshot overwrite data for LUNs and space-reserved files, to be used after all other space in the volume is used.

G

giveback The return of identity from an emulated storage system to the failed system, resulting in a return to normal operation. The reverse of *takeover*.

global namespace See *namespace*.

growth rate The measurement of how fast the storage is filling. The growth rate is determined by dividing the daily growth rate by the total amount of space in the storage system.

H

HA (high availability)

- In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an *HA pair*, that are configured to serve data for each other if one of the two nodes stops functioning.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

HA pair

- In Data ONTAP 8.x, a pair of nodes whose controllers are configured to serve data for each other if one of the two nodes stops functioning.

  Depending on the system model, both controllers can be in a single chassis, or one controller can be in one chassis and the other controller can be in a separate chassis.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

host A computer system that accesses data on a storage system.

host bus adapter (HBA) An interface card that plugs into a SAN device. SAN devices use the ports on their respective HBAs to connect to each other in a SAN. Each SAN device might contain one or more HBAs, and an HBA might contain more than one port. Each port can be used to establish a connection to a SAN.

I

igroup initiator group. A collection of unique iSCSI node names of initiators (hosts) in an IP network that are given access to *front-end LUNs* when they are mapped to those LUNs. (Array LUNs on a storage array that provide storage for V-Series systems can be considered *back-end LUNs*.)

incident An issue that has already impacted the availability or capacity of storage objects.

incremental transfer A subsequent backup after a baseline transfer has occurred of a primary directory in which only the new and changed data since the last backup (baseline or incremental) is transferred. The transfer time of incremental transfers can be significantly less than the baseline transfer.

Infinite Volume A logical entity contained in a Storage Virtual Machine (SVM, formerly known as Vserver)—referred to as SVM with Infinite Volume—that holds user data. An Infinite Volume obtains its storage from multiple aggregates.

initiator The system component that originates an I/O command over an I/O bus or network. The target is the component that receives this command.

inode A data structure containing information about files on a storage system and in a UNIX file system.

iSCSI Internet Small Computer Systems Interface (iSCSI) protocol. A licensed service on the storage system that enables you to export LUNs to hosts using the SCSI protocol over TCP/IP.

iSCSI router A storage router implementing the Internet Small Computer Systems Interface (iSCSI) protocol (SCSI over IP) to extend access of a Fibre Channel fabric and attached storage devices to IP servers.

J

JBOD Just a Bunch Of Disks. An array of disks without any redundancy; that is, without RAID configuration.

job A long-running operation, for example, scheduled local backup of a dataset, a mirror transfer, and password updates.

L

level-0 backup An initial backup (also known as a *baseline transfer*) of a primary directory to a secondary volume in which the entire contents of the primary directory are transferred.

LIF logical interface. Formerly known as *VIF* (virtual interface) in Data ONTAP GX. A logical network interface, representing a network access point to a node. LIFs currently correspond to IP addresses, but could be implemented by any interconnect. A LIF is generally bound to a physical network port; that is, an Ethernet port. LIFs can fail over to other physical ports (potentially on other nodes) based on policies interpreted by the LIF manager.

Lightweight Directory Access Protocol (LDAP) A client-server protocol for accessing a directory service.

local backup Local backup protection (also referred to as *Snapshot protection*) is the periodic capture of the active data on a NetApp storage system in backup images and the storage of those images on that same system. If active data on the local system is accidentally deleted or corrupted, it can quickly be restored with the most recent image stored locally from the last local backup job. Local backup operations are typically employed on the primary storage systems, where data is being actively updated and where, in event of accidental data loss, data restoration from the last hour or two might be required. Local backup protection is based on NetApp Snapshot technology.

local backup copy A copy of data, usually on a primary node, created using Snapshot technology and that resides on the primary dataset node.

Logical Unit Number (LUN) A SCSI identifier of a logical unit of storage on a target. LUNs are often referred to as *virtual disks*, and vice versa. See also *virtual disk*.

logical object The entity that represents a storage container, such as a volume, qtree, LUN, or dataset.

lower threshold The value set to generate an event when a counter falls and remains below that value for longer than the specified interval.

M

A (v) next to a term indicates that the definition is for the verb form of the word, while an (n) next to a term indicates that the definition is for the noun form of the word.

maintenance user The user who has access rights to deploy and configure an OnCommand Unified Manager virtual appliance.

Management Information Base (MIB) ASCII files that describe the information that the SNMP agent sends to network management stations.

member Any data object that subscribes to or is created by a storage service.

mirror (v) The process of creating an exact duplicate of all volume data from aNetApp source storage system to a destination storage system. If data in the source storage system is lost or made unavailable, then that replicated data can quickly be made available from the destination mirror site. Mirror operations are employed from primary to secondary storage and from secondary to tertiary storage. Mirror protection is based on NetApp Volume SnapMirror technology.

mirror copy (n) The exact duplicate of all volume data (both active and protected) from aNetApp source storage system to a destination storage system, created using NetApp Volume SnapMirror technology.

move (v) To physically move data and any needed associated configuration of an object from one aggregate to another within a cluster, including within a single node.

N

namespace In network-attached storage (NAS) cluster environments, an abstraction layer for data location that provides a single access point for all data in the system. It enables users to access data without specifying the physical location of the data, and enables administrators to manage distributed data storage as a single file system. Sometimes referred to as *global namespace*.

NDMP Network Data Management Protocol. A protocol that allows storage systems to communicate with backup applications and provides capabilities for controlling the robotics of multiple tape backup devices.

Network File System (NFS) export A service exposed from a NAS device to provide file-based storage through the NFS protocol. NFS is mostly used for UNIX-like operating systems, but other operating systems can access NFS exports as well.

node In Data ONTAP, a storage system or in a cluster or an HA pair.

To distinguish between the two nodes in an HA pair, one node is sometimes called the *local node* and the other node is sometimes called the *partner node* or *remote node*.

nondisruptive The ability of a system to continue serving data to clients during a system process or activity, such as a LUN restore operation or an online migration.

O

offline A database state indicating that the database is not available to users (for example, the database is in the following states: shutdown, started, or mounted).

online A database state indicating that the database is available to users (for example, open).

OnCommand administrator An RBAC role that enables a person to configure settings for items unrelated to storage management, such as user roles, security certificates, database access, LDAP, SMTP, networking, and AutoSupport.

operator An RBAC role that enables a person to view data and to view, assign, and resolve events in OnCommand Unified Manager.

P

parity disk The disk on which parity information is stored for a RAID4 disk drive array. In RAID groups using RAID-DP protection, two parity disks store the parity and double-parity information. Used to reconstruct data in failed disk blocks or on a failed disk.

partner node From the point of view of the local node (storage system), the other node in a high-availability configuration.

port A physical connection point on computers, switches, storage arrays, and so on, which is used to connect to other devices on a network. Ports on a Fibre Channel network are identified by their World Wide Port Name (WWPN) IDs. TCP/IP ports are used as virtual addresses assigned to each IP address.

policy A set of parameters that are grouped together as a distinct entity, so that the set of parameters can be applied to objects as a unit.

protection artifact An object, such as a destination data object, or a protection relationship that the Unified Manager server creates to support protection jobs when a data object is subscribed to a storage service.

protection policies The entities that enable you to set the automation controls for scheduling, monitoring, and alerts on any set of data in terms of normal backup, offsite backup, disaster and recovery backup, and regulatory copies.

protection relationship The SnapMirror or SnapVault relationship that exists between a source data object and a destination data object.

Q

qtree A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.

R

RAID-DP redundant array of independent disks, double-parity.

raw capacity The total amount of addressable blocks on physical disk drives. Calculated by multiplying the number of disk drives by the labeled capacity of those disk drives. For V-Series systems, it refers to the size of the array LUNs.

RBAC Role-based Access Control. A system whereby access to resources is decided based on the role of a user. RBAC controls who has access to various operations on which resources. Access to resources is first assigned to roles and roles are then assigned to users. Conforms to NIST RBAC standard.

recovery The re-creation of a past operational state of an entire application or computing environment. Compare *restore*. A recovery operation can encompass a restore. *Recovery* and *restore* are often used synonymously. Recovery and restore can also be *in-place*, meaning that copies are mounted and used locally instead of being copied elsewhere.

remote backup A copy of data on another set of physical disks or medium. Also referred to as *secondary storage*. See also *local backup*.

replication The process of duplicating data from one highly available site to another. The replication process can be synchronous or asynchronous. Duplicates are known as *clones*, *point-in-time copies*, or *Snapshot copies*, depending on the type of copy being made.

restore The copying of an object, such as a file or an attribute, or an entire application or virtual machine, back to its original source. Compare *recovery*. A restore can be part of a recovery operation. *Restore* and *recovery* are often used synonymously. Restore and recovery can also be *in-place*, meaning that copies are mounted and used locally instead of being copied elsewhere.

retention period The user-specified minimum length of time that a local backup Snapshot copy must be retained.

risk The issues that can impact the availability or capacity of storage objects.

root member	A data object that subscribes to a storage service.

S

SAN	storage area network. A dedicated network linking servers or workstations to devices, typically over Fibre Channel. SAN allows data and devices to be shared across a network as if they were attached locally.
severity type	The level of priority assigned to an event to help determine priorities for taking corrective action.
SFO	See *storage failover (SFO)*.
share	A directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known as a *CIFS share*.
Snapshot available capacity	The storage space that is available in the Snapshot reserve for Snapshot copies. Calculated by subtracting the total Snapshot used capacity from the Snapshot reserve capacity.
Snapshot copy	An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copies enable users to restore files and to back up the storage system to tape while the storage system is in use.
Snapshot overflow	The storage space that is consumed by Snapshot copies from the total data capacity of a volume or an aggregate. Calculated by subtracting the Snapshot reserve capacity from the Snapshot used capacity.
Snapshot reserve capacity	The storage space that is set aside by the volume or the aggregate for its Snapshot copies. Data cannot be written to this space.
Snapshot return	The capacity savings of Snapshot copies when compared to full volume copies. Calculated by the volume capacity - the Snapshot capacity.
Snapshot unused capacity	The Snapshot reserve space remaining after the Snapshot copies are created.
Snapshot used capacity	The storage space used by the Snapshot copies in a volume or an aggregate.
spare disk	A physical disk that is part of a storage device that is the same technology type (FC, SATA), size, and speed as a standard disk. A spare disk is used in case the standard disk malfunctions.
storable capacity	The disk capacity that is available for use after right-sizing.
standard HA configuration	A configuration set up so that one node automatically takes over for its partner when the partner node becomes impaired.

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

storage controller The component of a storage system that runs the operating system and controls its disk subsystem. Storage controllers are also sometimes called *controllers*, *storage appliances*, *appliances*, *storage engines*, *heads*, *CPU modules*, or *controller modules*.

storage efficiency The ratio of usable capacity to effective used capacity, accounting for efficiency returns. Calculated by the effective used capacity / the usable capacity.

storage failover (SFO) The method of ensuring data availability by transferring the data service of a failed node to another node in the cluster. Transfer of data service is often transparent to users and applications. Also referred to as *controller failover (CFO)*or *cluster failover (CFO)*.

storage utilization The ratio of usable capacity to used capacity, without accounting for efficiency returns. Calculated by the used capacity / the usable capacity.

Storage Virtual Machine (SVM) (Known as *Vserver* prior to clustered Data ONTAP 8.2.1. The term "Vserver" is still used in CLI displays and `vserver` command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs—*admin*, *node*, and *data*—but unless there is a specific need to identify the type of SVM, "SVM" usually refers to the data SVM.

SVM (Storage Virtual Machine; known as *Vserver* prior to clustered Data ONTAP 8.2.1. The term "Vserver" is still used in CLI displays and `vserver` command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs—*admin*, *node*, and *data*—but unless there is a specific need to identify the type of SVM, "SVM" usually refers to the data SVM.

SVM guaranteed available capacity The data storage space that is guaranteed by the SVM to its underlying volumes but is not used. Calculated as the sum of the available size of all the thick provisioned volumes.

SVM used capacity The data storage space that is guaranteed by the SVM to its underlying volumes. Calculated as the sum of the used data capacity of all the volumes associated with the SVM.

SVM unguaranteed capacity The data storage space that is not guaranteed by the SVM to its underlying volumes. Calculated as the sum of the available sizes of all the thin provisioned volumes.

SVM total capacity The sum of the total data storage space of all the volumes in the SVM.

SVM unguaranteed capacity	The data storage space that is not guaranteed by the SVM to its underlying volumes. Calculated as the SVM total capacity minus the sum of the committed capacity of aggregates that are associated with the SVM.
system reserve capacity	The capacity required for fixed system reserves, RAID parity, mirroring, and spare drives. Calculated by the fixed reserve + RAID reserve + spares.

T

takeover	The emulation of the failed node identity by the takeover node in a high-availability configuration; the opposite of *giveback.*
takeover node	A node (storage system) that remains in operation after the other node stops working and that hosts a virtual node that manages access to the failed node disk shelves and network connections. The takeover node maintains its own identity and the virtual node maintains the failed node identity.
thin provisioning (TP)	A method of optimizing the efficiency with which the available space is used in storage. When many applications share access to the same storage array, thin provisioning enables administrators to maintain a single free space buffer pool to service the data requirements of all applications. This is done by allocating disk storage space in a flexible manner among multiple consumers, based on the minimum space required by each at any given time. This avoids poor utilization that occurs on traditional storage arrays where large pools of storage capacity are allocated to individual applications, but much remains unused.
trap	An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the storage system.

U

unassigned disk	A disk that is not assigned to any node or counted as a spare disk.
unused capacity	The free, usable storage space available for storing user data on the device, excluding capacity reserved for Snapshot copies or data.
unused reserve capacity	The capacity allocated but unused by Aggregate Snapshot Reserve, Volume Snapshot Reserve, Volume Fractional Reserve, and Vol/LUN/File Guaranteed Space. These reserves are adjustable by the user. Calculated by the aggregate snapshot unused reserve + volume snapshot unused reserve + volume fractional unused reserve + vol/LUN/file unused guaranteed space.
usable capacity	The capacity available to applications and users. Calculated by the raw capacity - the system reserve capacity.

used capacity The capacity used by application or user data, including volume Snapshot copies and aggregate Snapshot copies. Calculated by the usable capacity - the free capacity.

V

vApp See *virtual appliance*.

virtual appliance A prebuilt software solution containing virtual machines and software applications that are integrated, managed, and updated as a package. Also called *vApp*.

virtual machine A guest operating system and any application installed thereon, running on a computing device on which the software is installed, or suspended to disk or any other storage media accessible by the computing device.

VMware VirtualCenter (VC) A management software suite used to create your VMware datastores and virtual machines and to configure the storage system volumes as the containers in which your active datastore and virtual machine images reside. It consists of the following components:

- VMware agents—software modules installed on an ESX server to carry out VC server requests
- VirtualCenter (VC) server—a server communicating with VMware agents on an ESX server
- Virtual Infrastructure (VI) client—a GUI client to manage the VC server
- VMware ESX server—an enterprise-level product that integrates server processes, storage functionality, and networking resources into multiple virtual systems.

  ESX server can also refer to a physical host running an ESX server hypervisor OS.

volume

- For Data ONTAP, a logical entity that holds user data that is accessible through one or more of the supported access protocols, including Network File System (NFS), Common Internet File System (CIFS), Fibre Channel (FC), and Internet SCSI (iSCSI). V-Series treats an IBM volume as a disk.
- For IBM, the area on the storage array that is available for a V-Series system or non V-Series host to read data from or write data to. The V-Series documentation uses the term *array LUN* to describe this area.

volume committed capacity The data storage space in a volume that is committed to provide storage space for its underlying qtrees based on their quota settings. Calculated as the sum of all qtrees disk hard limits. The sum does not include the qtrees for which the disk hard limit is not set.

volume total capacity	The data storage space in a volume that can be used by qtrees, LUNs, or other files and volume-level Snapshot copies. Calculated as the total data capacity of the volume plus the volume-level Snapshot reserve space.
Vserver	(Known as “Storage Virtual Machine (SVM)” in clustered Data ONTAP 8.2.1 and later.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of Vservers—*admin*, *node*, and *cluster* (“cluster Vserver” is called “data Vserver” in Data ONTAP 8.2)—but unless there is a specific need to identify the type of Vserver, “Vserver” usually refers to the cluster/data Vserver.

Copyright information

Copyright © 1994–2014 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at *www.ibm.com/legal/copytrade.shtml*.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

about
    snapshot copy reserve *138*
access checks
    introduction to using RBAC to enable application administrator *320*
access roles (RBAC)
    See RBAC
acknowledging
    events *54*
Active Directory
    setting up authentication services *328*
Add Alert dialog box *93*
Add Cluster dialog box *155*
Add User dialog box *324*
adding
    a condition or condition set *303*
    alerts *77*
    authentication servers *329*
    clusters *120*
    new rules *299*
    notes about an event *53*
    OnCommand Workflow Automation settings *292*
    rules *298*
administrator roles
    See RBAC
administrators
    introduction to using RBAC to restrict functionality
    access to selected *320*
    OnCommand *321*
    storage *321*
Aggregate details page *178*
aggregate events
    impact levels *57*
    listed by impact area *57*
    severities *57*
    source types *57*
Aggregate Thresholds Setup Options dialog box *102*
aggregates
    adding to resource pools *253*, *255*
    configuring global threshold values for *98*
    details about *178*
    editing threshold settings *101*
    exporting data to .csv files *18*
    global thresholds for *102*
    removing from resource pools *243*, *253*
    viewing annotations for *209*
    viewing inventory list and details *127*
Aggregates dialog box *255*
Aggregates page *175*
alerts
    adding *77*
    configuring your environment for *24*
    creating *77*
    definition *77*
    deleting *81*
    editing *81*
    guidelines for adding *79*
    modifying *81*
    testing *80*
    viewing details *80*
annotations
    creating rules for storage objects *208*
    definition of *208*
    description of types *210*
    for volumes, details *190*
    managing *210*
    removing *209*
    rules for *212*
    using for storage objects *208*
    viewing *210*
    viewing for storage objects *209*
assigning
    events *52*
associations
    creating SVM *245*
    evaluating sufficiency of SVM *247*
    viewing SVM *247*
authentication
    Active Directory *332*
    adding servers *329*
    deleting servers *331*
    editing servers *331*
    enabling remote *327*
    OpenLDAP *332*
    servers *332*
    setup options *28*, *332*
    testing for remote users and groups *330*
    testing remote users and groups *317*
    troubleshooting unknown authentication error *344*
authentication services

setting up using Active Directory *328*
setting up using OpenLDAP *328*
autogrow
enabling on destination *221*
reasons for enabling *261*
AutoSupport
enabling periodic messages *25*
sending on-demand messages *344*
setup options *28*
AutoSupport Setup Options dialog box
fields defined *32*
availability
details about clusters *35*
details in Quick Takes dashboard area *35*
details in Unresolved Incidents and Risks dashboard area *36*

B

benefits
of Unified Manager *14*
bookmarks
for favorite topics *18*
Browse Directories dialog box *265*

C

capabilities
table of roles associated with *322*
capacity
details in Quick Takes dashboard area *35*
details in Unresolved Incidents and Risks dashboard area *36*
information for volumes *190*
capacity events
node root aggregates *119*
certificates
downloading signing requests for HTTPS security *340*
generating HTTPS security certificates *338*
installing HTTPS security certificates *341*
security, downloading signing request *342*
security, regenerating *342*
security, viewing *342*
viewing HTTPS security *337*
CIFS
about *143*
license *143*
shares inaccessible, troubleshooting access issues *347*
viewing the share details *129*
CIFS (Common Internet File System) *143*
clients
SMB, supported *143*
cluster
description of *132*
Cluster availability details pane *35*
cluster components
definition *45*
limits imposed on workloads by *45*
Cluster details page *146*
cluster events
impact levels *59*
listed by impact area *59*
severities *59*
source types *59*
cluster monitoring
understanding *118*
cluster network *138*
clustered Data ONTAP systems
See clusters
clusters
adding *120*
availability details *35*
creating rules to annotate *208*
deleting *123*
details about *146*
editing settings *122*
exporting data to .csv files *18*
how monitoring works *118*
how the discovery process works *20*
rediscovering if discovery fails *122*
rediscovering manually *123*
removing *123*
removing if discovery fails *121*
viewing discovery status *120*
viewing inventory list and details *124*
clusters and HA pairs *138*
Clusters page *144*
Common Internet File System (CIFS) *143*
compression
enabling on destination *221*
reasons for enabling *261*
condition sets
defined *309*
conditions
defined *309*
Configure Network Settings dialog box *31*
Configure Protection wizard
purpose *272*

tasks enabled using *272*
Configure Workflow Annotation Rules dialog box *212*
configuring
aggregate global threshold values *98*
DNS *25*
environment options *25*
event settings *48*
global options *26*
information for volumes *190*
network settings *25*
notification settings *49*
thresholds *100*
volume global threshold values *99*
your environment *22*
connections
between Performance Manager and Unified Manager, purpose of *44*
constituents
for infinite Volumes *298*
viewing *294*
Create Resource Pool dialog box
purpose *240*
Create Rule dialog box
using to create rules for data policies *310*
Create Schedule dialog box
purpose *278*
Create SnapMirror Policy dialog box
purpose *279*
Create SnapVault dialog box
purpose *280*
Create SVM Associations wizard
about *258*
tasks enabled *258*
creating
alerts *77*
custom rules *300*
resource pools *240*
rules *298*
rules to generate email addresses *110*
rules using templates *299*
SVM associations *245*, *257*
creating email notification formats
for user and user group quotas *111*
custom rules
creating *300*
editing directory paths *303*
editing file types *303*
editing owners *303*
customization
of environment options *25*
of window layouts *16*

D

dashboard
description of areas and panes *33*
monitoring from *33*
Quick Takes area *35*
understanding *33*
Unresolved Incidents and Risks area *36*
using to view details of performance incidents *41*
using to view incidents that affect performance *41*
data
exporting to .csv files *18*
restoring from the Volume details page *239*
restoring from the Volumes page *238*
restoring to a directory other than the original source *265*
data constituents
definition *298*
data network *138*
data policies
configuration details *160*
defined *307*
deleting rules *304*
exporting *305*
how the default rule works *308*
how they filter files for SVMs with Infinite Volume *308*
importing *306*
previewing changes to the configuration *304*
understanding *307*
data protection mirror copies
namespace constituent in Infinite Volumes, of *298*
data protection transfers
temporarily halting *227*
data restoration
to a directory other than the original source *265*
data sources
description of *205*
viewing details about *129*
Data Sources page *205*
database users
creating *316*
defined *321*
deduplication
enabling on destination *221*
reasons for enabling *261*
default rules
defined *308*

definitions
performance incidents *43*
deleting
a cluster if discovery fails *121*
alerts *81*
annotations from storage objects *209*
authentication servers *331*
clusters *123*
connection with a Performance Manager server *40*
OnCommand Workflow Automation settings *292*
OnCommand Workflow Automation setup *291*
resource pools *243*
rules *304*
Snapshot copies *130*
SVM associations *247*, *257*
user group quota or user quota email addresses *111*
users *318*
destination and source associations
creating *245*
destination volume efficiency
space guarantee *262*
DHCP
enabling *25*
Disable Events dialog box *91*
disabling
events *51*, *90*, *91*
discovery
of clusters, how the process works *20*
of clusters, manual *123*
restarting cluster discovery *122*
viewing status of cluster *120*
disk space
viewing the list of Snapshot copies for calculating *130*
disks events
impact levels *60*
listed by impact area *60*
severities *60*
source types *60*
display
minimum recommended setting *15*
DNS
configuring *25*
domain controllers
SMB, supported *143*

E

Edit Aggregate Thresholds dialog box *268*
Edit Alert dialog box *95*
Edit Cluster dialog box *156*
Edit Relationship dialog box
purpose *281*
Edit Resource Pool dialog box *254*
Edit Rule dialog box *311*, *312*
Edit Storage Class Thresholds dialog box *314*
Edit SVM with Infinite Volume dialog box *313*
Edit User dialog box *325*
Edit Volume Thresholds dialog box *270*
editing
aggregate threshold settings *101*
alerts *81*
authentication servers *331*
cluster settings *122*
global options *26*
Infinite Volume threshold settings *295*
network settings *25*
resource pools *241*, *254*
storage class threshold settings *296*
unmanaged relationship lag threshold settings *99*
user group quota or user quota email addresses *111*
user settings *317*
volume threshold settings *102*
efficiency
information for volumes *190*
space on secondary volume *261*
efficiency settings
enabling on destination *221*
Email Notification Format dialog box *114*
enabling
autogrow, why you do it *261*
compression, why you do it *261*
deduplication, why you do it *261*
DHCP *25*
events *51*, *90*
periodic AutoSupport messages *25*
enclosures events
impact levels *61*
listed by impact area *61*
severities *61*
source types *61*
environment
configuration options *25*
setup *22*
Event details page *87*
event impact areas
availability *56*
capacity *56*
configuration *56*
description *56*

performance *56*
protection *56*
event impact levels
description *56*
event *56*
incident *56*
risk *56*
event publishers
defined *321*
Event Retention Settings *82*
event severity types
critical *55*
description *55*
error *55*
information *55*
warning *55*
event states
definition of *55*
events
acknowledged event, definition *55*
acknowledging *54*
adding notes about *53*
aggregate *57*
assigning to users *52*
cluster ports *60*
clusters *59*
configuring settings for *48*
definition of *48*
details *87*
disabling *51*, *91*
disks *60*
enabling *51*
enclosures *61*
exporting data to .csv files *18*
fans *62*
flash card *62*
generated by performance incidents *43*
impact areas *56*
impact levels *56*
inodes *62*
list of *57*
logical interface *63*
LUN *63*
managing *90*
new event, definition *55*
nodes *64*
NVRAM battery *65*
obsolete event, definition *55*
ports *60*
power supplies *65*
protection *66*
qtree *66*
resolved event, definition *55*
resolving *54*
reviewing notes about *53*
service processor *67*
setup options *82*
severity and object status *57*
severity information *85*
severity types *55*
SnapMirror relationships *67*
SnapVault relationship *68*
states defined *55*
status *85*
storage failover settings *69*
storage services *69*
storage shelf *70*
SVM *74*
SVM storage class *76*
user and group quota *70*
viewing details *50*
viewing notes about *53*
viewing unassigned *53*
volume *71*
volume move status *73*
Events page
fields in the *85*
Export Policy Rules dialog box *204*
exporting
data to .csv files *18*

F

failed cluster discovery
removing *121*
failed protection jobs
retrying *250*
fans events
impact levels *62*
listed by impact area *62*
severities *62*
source types *62*
favorites
adding topics to bookmarks list *18*
features
in Unified Manager *13*
file system
data policy, impact on *308*
files
maximum number on Infinite Volume *297*

flash card events
impact levels *62*
listed by impact area *62*
severities *62*
source types *62*
FlexVol volumes
with SVMs, explained *132*

G

general settings
setup options *28*
giveback
definition of *142*
global options
editing *26*
group quota events
impact levels of *70*
groups
introduction to using RBAC to define user roles for managing *320*
testing remote authentication *317*, *330*
viewing quota details *109*
guidelines
for adding alerts *79*

H

HA configurations
definition of *138*
HA pairs and clusters *138*
hard quotas
how they work *113*
hardware
components described *142*
HA components described *142*
Help
using *17*
high annotation types
definition of *210*
host names
changing *23*
HTTPS
changing the protocol to *122*
downloading security certificate signing requests *340*
setup options *342*
viewing the security certificate *337*
HTTPS certificates
downloading a new *27*
generating new security certificates *338*
installing a new *27*
installing security certificates *341*
regenerating *27*
troubleshooting issue with installing or regenerating on a Unified Manager server enabled for performance monitoring *348*
viewing *27*
working with *27*

I

I/O performance incidents
defined *43*
icons misaligned
troubleshooting *345*
impact areas
availability *56*
capacity *56*
configuration *56*
description *56*
performance *56*
protection *56*
impact levels
description *56*
event *56*
incident *56*
risk *56*
incidents
I/O performance, defined *43*
unresolved, dashboard details for *36*
Infinite Volumes
creating custom rules for *300*
creating rules for placing data in *299*
data constituents *298*
definition of *297*
editing threshold settings *295*
maximum number of files *297*
monitoring
Infinite Volumes *294*
namespace constituent *298*
namespace mirror constituents, definition of *298*
previewing changes to data policies of *304*
storage classes, definition of *297*
SVMs with Infinite Volume
viewing details about *294*
viewing constituents of *294*
viewing details about *294*
with SVMs, explained *132*
initialization operations

default Snapshot copy behavior for *285*
selecting maximum transfer rate for *231*, *232*
selecting Snapshot copies for *231*, *232*
selecting transfer priority for *231*, *232*
Initialize/Update dialog box
purpose *282*
inodes events
impact levels *62*
listed by impact area *62*
severities *62*
source types *62*
introduction
to Unified Manager *13*

J

Job Details page
purpose *266*
jobs
aborting *249*, *259*
defined *248*
list of those you can monitor *248*
monitoring *248*
progress monitoring *266*
refreshing *259*
retrying *259*
retrying failed protection *250*
status *266*
stopping *249*
troubleshooting failures *266*
viewing information about *259*
viewing tasks associated with *249*
Jobs page
purpose *259*

L

LDAP server slow to respond
troubleshooting *346*
LDAP user of OpenLDAP server
troubleshooting *346*
licenses
CIFS *143*
LIF events
See logical interface events
LIFs
about *136*
local users
changing password for *319*
creating *316*
defined *321*
logical interface events
impact levels *63*
listed by impact area *63*
severities *63*
source types *63*
logical storage
defined *135*
low annotation types
definition of *210*
LUN events
impact levels *63*
listed by impact area *63*
severities *63*
source types *63*

M

maintenance console
accessing using Secure Shell *339*
accessing using VM console *340*
restarting the virtual machine *338*
restarting Unified Manager *338*
role of maintenance user *320*
maintenance user
defined *321*
what this user does *320*
Manage Alerts page *92*
Manage Annotations page *210*
Manage Events page
events list *90*
Manage users page *323*
management server
authentication setup options *28*
AutoSupport setup options *28*
HTTPS setup options *28*
managing
annotations *210*
events *90*
SVM associations *257*
users *323*
manual rediscovery
of clusters *123*
maximum transfer rates
changing using Edit Relationship *281*
editing for existing relationships *222*, *223*
when initializing relationships *282*
when updating relationships *282*
messages
sending on-demand AutoSupport *344*

misalignment of icons
troubleshooting *345*
mission-critical annotation types
definition of *210*
modifying
aggregate threshold settings *101*
Infinite Volume threshold settings *295*
OnCommand Workflow Automation settings *292*
storage class threshold settings *296*
unmanaged relationship lag threshold settings *99*
user settings *317*
volume threshold settings *102*

N

namespace
data policy, impact on *308*
namespace constituent
and maximum files *297*
namespace constituents
definition *298*
namespace mirror constituents
definition of *298*
navigation
understanding window layouts *15*
nested groups
disabling support for to speed performance *346*
network settings
configuring *25*
customizing the host name *23*
editing *25*
node events
impact levels *64*
listed by impact area *64*
severities *64*
source types *64*
node root volumes
monitoring *119*
nodes
exporting data to .csv files *18*
single node cluster
See single-node cluster
viewing inventory list and details *125*
Nodes page
information displayed and tasks enabled in *156*
notification
adding alerts *77*
configuring settings *49*
setup options *83*
Notification Setup Options dialog box *83*
NVRAM events
impact levels *65*
listed by impact area *65*
severities *65*
source types *65*

O

objects
status as related to event severity *57*
on-demand AutoSupport messages
sending *344*
OnCommand administrators
defined *321*
OnCommand Workflow Automation
pairing with Unified Manager *290*
removing *291*
setup options *292*
online Help
using *17*
OpenLDAP
setting up authentication services *328*
operators
defined *321*

P

pairing
Workflow Automation *290*
partner applications
ensuring conformance of *245*
SVM association and resource pool requirements *245*
passwords
changing local user *319*
performance incidents
definition of *43*
discovery and analysis of *47*
display of, in Unified Manager *44*
generation on a cluster component *45*
viewing all *42*
viewing details of *41*
viewing details of incidents listed on the Events page *42*
performance issues
viewing incidents associated with *41*
Performance Manager
about connections between multiple Performance Manager servers, and Unified Manager *44*

configuring a connection to a Unified Manager server *39*
deleting a connection to a Unified Manager server *40*
Maintenance console
to delete a performance monitoring connection with a Unified Manager server *40*
purpose of connection with Unified Manager *44*
performance monitoring
configuring connections between Performance Manager and Unified Manager *39*
deleting connections between Performance Manager and Unified Manager *40*
disabling *40*
enabling *39*
introduction to enabling *38*
troubleshooting disablement due to HTTPS certificate installation or regeneration *348*
understanding key concepts *43*
periodic AutoSupport messages
enabling *25*
physical storage
adding clusters *120*
Clusters page *144*
editing cluster settings *122*
Nodes page *156*
removing clusters *123*
policies
creating for SnapMirror relationships *224*
creating for SnapVault relationships *224*
exporting data policies *305*
importing data policies *306*
port events
impact levels *60*
listed by impact area *60*
severities *60*
source types *60*
ports
concepts *137*
editing, for authentication servers *331*
power supplies events
impact levels *65*
listed by impact area *65*
severities *65*
source types *65*
protection
details in Quick Takes dashboard area *35*
details in Unresolved Incidents and Risks dashboard area *36*
information for volumes *190*
monitoring jobs for *248*
viewing inventory list *215*
protection events
impact levels *66*
listed by impact area *66*
severities *66*
source types *66*
protection jobs
retrying failed *250*
protection policies
changing using Edit Relationship *281*
editing for existing relationships *222*, *223*
protection relationships
aborting transfers for *225*, *226*
creating from the Volume details page *217*, *218*
creating from the Volumes page *216*, *219*
editing for existing *222*, *223*
editing from the Volume details page *223*
editing from the Volumes page *222*
initializing from the Volume details page *231*, *232*
quiescing from the Volume details page *227*
removing from the Volume details page *229*
reversing *236*
reversing from the Volume details page *235*
updating from the Volume details page *231*, *232*
protection schedules
advanced *278*
basic *278*
changing using Edit Relationship *281*
editing for existing relationships *222*, *223*
Protection view *184*
provisioning
adding members to resource pools *242*
deleting secondary relationships between source and destination SVMs *247*
destination volumes, creating SVM associations *245*
secondary volumes, increasing space to resource pools for *242*

Q

qtree events
impact levels *66*
listed by impact area *66*
severities *66*
source types *66*
qtrees
creating rules for *110*
identifying inaccessible *348*
unable to access or write to *348*
Quick Takes area *35*

quiesced protection transfers
resuming from the Volume details page *230*
resuming from the Volumes page *229*
quotas
creating email notification formats for users and user groups *111*
events about user and group *70*
group *113*
how they work *113*
qtree *113*
user *113*
viewing details of user or group *109*
volume *113*
why you use *113*

R

RBAC
definition *320*
introduction to managing groups of users by using *320*
reassigning
events *52*
rediscovery
cluster *122*
of clusters, manual *123*
relationships
deleting secondary provisioning between source and destination SVMs *247*
unmanaged, editing lag thresholds settings for *99*
unmanaged, global thresholds for *267*
Relationships Thresholds Setup Options dialog box
purpose *267*
remote authentication
enabling *327*
servers *332*
setup options *332*
remote groups
adding *316*
defined *321*
testing authentication *317*, *330*
remote users
adding *316*
defined *321*
testing authentication *317*, *330*
removing
alerts *81*
clusters *123*
failed cluster discovery *121*
OnCommand Workflow Automation setup *291*
replication labels
creating for SnapVault policies *280*
requirements
when you upgrade or reinstall OnCommand Workflow Automation *291*
resolving
events *54*
resource pool members
adding *251*
deleting *251*
displaying *251*
resource pools
about *215*
adding aggregates to *253*, *255*
adding members to *242*, *251*
configuring to support cascaded protection in storage services *245*
creating *240*, *251*
deleting *243*, *251*
deleting members from *251*
describing *253*
editing *241*, *251*, *254*
managing *251*
monitoring *251*
monitoring available capacity *241*
naming *253*
removing aggregates from *253*
removing member aggregates from *243*
renaming *241*
viewing inventory *241*
Resource Pools page *251*
resources
selecting using SVM associations *244*
Restore dialog box
purpose *263*
restoring data
to a directory other than the original source *265*
resynchronization operations
default Snapshot copy behavior for *285*
selecting maximum transfer rate for *233*, *234*, *284*
selecting Snapshot copies for *233*, *234*
selecting transfer priority for *233*, *234*, *284*
Resynchronize dialog box
purpose *284*
retention count
changing for SnapVault policies *280*
Reverse Resync dialog box
purpose *288*
reverse resync operation
performing from the Volumes page *236*

reverse resync operations
performing from the Volume details page *235*
reviewing
notes about events *53*
risks
unresolved, dashboard details for *36*
role-based access control
See RBAC
roles
assigning to users *316*
creating a user with the Event Publisher user *38*
defined *321*
table of capabilities associated with *322*
root aggregates
capacity events *119*
thresholds *119*
root volumes
monitoring *119*
rules
adding *298*
creating *110*
creating to annotate storage objects *208*
creating, custom *298*, *300*
creating, using templates *298*
default, defined *308*
defined *307*
deleting *304*
editing custom rules *303*
editing properties
directory paths *302*
file types *302*
owner *302*
editing template-based rules *302*
exporting *305*
how annotation rules work *212*
how they work *115*
importing *306*
previewing changes to *304*
templates, defined *309*
understanding *307*
viewing list of *301*
Rules to Generate User and Group Quota Email Address dialog box *115*
rules using templates
creating *299*

S

scheduled transfers
resuming from the Volume details page *230*
resuming from the Volumes page *229*
schedules
creating for SnapMirror relationships *278*
creating for SnapMirror transfers *221*
creating for SnapVault relationships *278*
creating for SnapVault transfers *221*
search bar
using to find storage objects *21*
searches
information about product features *17*
searching
for storage objects *21*
secondary volume efficiency
advanced settings *261*
secondary volume provisioning
increasing space to resource pools for *242*
Secure Shell
using to access the maintenance console *339*
security certificates
downloading a new *27*
downloading signing request *342*
downloading signing requests for HTTPS *340*
generating, HTTPS *338*
installing a new *27*
installing, HTTPS *341*
regenerating *27*, *342*
viewing *27*, *342*
viewing HTTPS *337*
working with *27*
Select Source Snapshot Copy dialog box
purpose *285*
servers
about connections between multiple Performance Manager servers and Unified Manager *44*
service processor events
impact levels *67*
listed by impact area *67*
severities *67*
source types *67*
Set Up OnCommand Workflow Automation dialog box *292*
setting up
aggregate global threshold values *98*
event settings *48*
notification settings *49*
OnCommand Workflow Automation *292*
SMTP server *49*
SNMP *49*
thresholds *100*
volume global threshold values *99*

setup
authentication options *332*
editing global options *26*
enabling periodic AutoSupport messages *25*
HTTPS options *342*
post-deployment *22*
setup options
sending on-demand AutoSupport messages *344*
Setup Options dialog box
general settings *28*
management server options *28*
threshold options *28*
severity types
critical *55*
description *55*
error *55*
information *55*
list of *57*
warning *55*
shares
viewing details about CIFS *129*
single-node cluster
description of *132*
SMB
supported clients *143*
supported domain controllers *143*
SMTP
options *83*
SnapDiff
namespace mirror constituents for *298*
SnapMirror events
impact levels *67*
listed by impact area *67*
severities *67*
source types *67*
SnapMirror policies
editing *222*, *223*
SnapMirror relationships
breaking *228*
creating from the Volume details page *218*
creating from the Volumes page *219*
creating policies for *224*
creating schedules for *278*
SnapMirror schedules
editing *222*, *223*
Snapshot copies
calculating reclaimable space *131*
deleting *130*
selecting for initialization operations *231*, *232*
selecting for resynchronization operations *233*, *234*, *284*
selecting for update operations *231*, *232*
selecting initialization default *282*
selecting relationship update default *282*
understanding *144*
viewing list of *130*
Snapshot copies on a Volume dialog box *206*
SnapVault policies
creating *224*
creating replication labels for *280*
editing *222*, *223*
editing retention count for *280*
setting transfer priorities for *280*
SnapVault relationship events
impact levels *68*
listed by impact area *68*
severities *68*
source types *68*
SnapVault relationships
creating *217*
creating from the Volumes page *216*
creating policies for *224*
creating schedules for *278*
SnapVault schedules
editing *222*, *223*
SNMP
options *83*
versions supported *83*
soft quotas
how they work *113*
source and destination associations
creating *245*
space guarantee
enabling on destination *221*
settings on destination *262*
space management
calculating reclaimable space for Snapshot copies *131*
storage administrators
defined *321*
storage capacity thresholds
definition of *98*
storage classes
capacity details of *160*
definition of *297*
editing threshold settings *296*
how data policies work *308*
storage failover events
impact levels *69*

listed by impact area *69*
severities *69*
source types *69*
storage objects
creating rules to annotate *208*
removing annotations *209*
searching for *21*
viewing annotations for *209*
storage performance health
viewing all performance incidents to assess *42*
storage services
SVM association and resource pool requirements *245*
storage services events
impact levels *69*
listed by impact area *69*
severities *69*
source types *69*
storage shelf events
impact levels *70*
listed by impact area *70*
severities *70*
source types *70*
Storage Virtual Machine details page *160*
SVM associations
about *244*
configuring to support cascaded protection in storage services *245*
creating *245*, *257*
deleting *247*, *257*
evaluating sufficiency of *247*
list of *257*
viewing *247*
SVM Associations page
content displayed *257*
tasks enabled *257*
SVMs
benefits *134*
creating rules for *110*
creating rules to annotate *208*
details about *160*
events listed by impact area *74*
exporting data to .csv files *18*
impact levels
of events *76*
severities
of events *76*
source types
of events *76*
storage class events listed by impact area *76*
viewing annotations for *209*
viewing inventory list and details *126*
with FlexVol volumes, explained *132*
with Infinite Volume, explained *132*
SVMs page *158*
SVMs with FlexVol volumes
explained *132*
SVMs with Infinite Volume
explained *132*
how a data policy filters file in *308*
previewing changes to data policies *304*
viewing constituents of *294*
syntax
for rules *115*, *212*

T

takeover
automatic *141*
definition of *142*
manual *141*
when it occurs *141*
tasks
viewing details about *249*
viewing information about *266*
template-based rules
editing *302*
templates
for rules *309*
testing
alerts *80*
authentication for remote users and groups *330*
remote users and groups *317*
thresholds
configuring *100*
description of *109*
editing settings for aggregates *101*
editing settings for Infinite Volumes *295*
editing settings for storage classes *296*
editing settings for unmanaged relationships *99*
editing settings for volumes *102*
for aggregates, global *102*
for unmanaged relationships, global *267*
for volumes, global *105*
global values for aggregates *98*
global values for volumes *99*
root aggregates *119*
setup options *28*
storage capacity, definition *98*
user quota, definition of *109*

time
how timestamps work in Unified Manager *21*
time zones *21*
timestamps
how they work in Unified Manager *21*
transfer checkpoints
clearing during an abort operation *225*, *226*
transfer priorities
for relationship updates *282*
setting for SnapMirror policy *279*
specifying for SnapMirror relationships *224*
when initializing relationships *282*
transfers
creating schedules for data protection *221*
resuming on quiesced protection relationship *229*, *230*
troubleshooting
access issues with volumes or qtrees *348*
adding LDAP user of OpenLDAP server *346*
CIFS share access issues *347*
job failures *266*
LDAP server slow to respond *346*
misalignment of icons *345*
unknown authentication error *344*
user not found *345*
types
of annotations *210*
of users *321*

U

Unified Manager
capabilities of *14*
introduction to *13*
Unified Manager servers
troubleshooting issue with installing or regenerating HTTPS certificates on when performance monitoring is enabled *348*
unknown authentication error
troubleshooting *344*
unmanaged relationships
editing lag thresholds settings for *99*
global thresholds for *267*
unresolved incidents
dashboard details for *36*
Unresolved Incidents and Risks area
details about *36*
unresolved risks
dashboard details for *36*
update operations
default Snapshot copy behavior for *285*
selecting maximum transfer rate for *231*, *232*
selecting Snapshot copies for *231*, *232*
selecting transfer priority for *231*, *232*
user group quotas
deleting *111*
editing *111*
user groups
creating email notification formats for quotas *111*
creating rules for *110*
user not found error
troubleshooting *345*
user quota events
impact levels of *70*
user quota thresholds
definition of *109*
user quotas
deleting *111*
editing *111*
user roles
assigning *316*
users
adding *316*
capabilities associated with *322*
changing password for local *319*
creating *316*
creating a user with the Event Publisher role *38*
creating email notification formats for quotas *111*
creating rules for *110*
deleting *318*
editing settings *317*
introduction to using RBAC to manage groups of *320*
maintenance user *320*
managing *323*
modifying settings *317*
roles *321*
testing remote authentication *317*, *330*
troubleshooting LDAP user of OpenLDAP server *346*
types *321*
viewing *318*, *323*
viewing quota details *109*
Users and Roles capability
See RBAC

V

viewing
aggregates list and details *127*

alert details *80*
annotation list *210*
CIFS share details *129*
clusters list and details *124*
constituents *294*
data sources *129*
disabled events list *90*
discovery status of clusters *120*
event details *50*
nodes list and details *125*
notes about events *53*
protection inventory list *215*
rules list *301*
SVM associations *247*
SVM associations list *257*
SVM list and details *126*
unassigned events *53*
user and group quotas *109*
users *318*
users list *323*
volumes list and details *128*
virtual machine console
accessing the maintenance console *340*
virtual machines
restarting *338*
VM console
accessing the maintenance console *340*
Volume details page
aborting SnapMirror relationships from *228*
aborting transfers from *225*
creating SnapMirror relationships using *218*
creating SnapVault relationships using *217*
editing protection relationships from *223*
initializing relationships from *231*, *232*
performing a reverse resync from *235*
quiescing protection relationships from *227*
removing relationships from *229*
restoring data from *239*
resynchronizing relationships from *233*
updating relationships from *231*, *232*
using to resume scheduled transfers *230*
volume events
impact levels *71*
listed by impact area *71*
severities *71*
source types *71*
volume move events
impact levels *73*
listed by impact area *73*
severities *73*
source types *73*
Volume Thresholds Setup Options dialog box *105*
volumes
capacity information *190*
configuration information *190*
configuring global threshold values for *99*
creating rules for *110*
creating rules to annotate *208*
details about *190*
editing threshold settings *102*
efficiency information *190*
exporting data to .csv files *18*
global thresholds for *105*
how they work *135*
identifying inaccessible *348*
Infinite Volume defined *297*
node root, monitoring *119*
Overview *184*
protection information *190*
Protection view *184*
provisioning using SVM associations *244*
restoring data *263*
Snapshot copies
understanding *144*
SVMs with FlexVol, explained *132*
SVMs with Infinite, explained *132*
unable to access or write to *348*
viewing annotations for *209*
viewing inventory list and details *128*
viewing the list of Snapshot copies *130*
Volumes page
aborting transfers from *226*
creating SnapMirror relationships from *219*
creating SnapVault relationships from *216*
editing protection relationships from *222*
performing a reverse resync operation from *236*
restoring data from *238*
resynchronizing relationships from *234*
using to resume scheduled transfers *229*
Vservers
See SVMs

W

window layouts
customization of *16*
samples of typical *15*
Workflow Automation
pairing with Unified Manager *290*
retrying jobs for *259*

what happens when you upgrade or reinstall *291*
workloads
definition *45*
system *45*
user-defined *45*