




# OnCommand® Unified Manager 6.0

## Administration Guide



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](http://www.netapp.com)  
Feedback: [doccomments@netapp.com](mailto:doccomments@netapp.com)

Part number: 215-08057\_A0  
July 2013



# Contents

<b>Unified Manager product and Administration Guide overview .....</b>	<b>6</b>
<b>Concepts related to working with Unified Manager .....</b>	<b>7</b>
What a cluster is .....	7
What Vservers are .....	8
How volumes work .....	9
What a FlexVol volume is .....	9
Capabilities that FlexVol volumes provide .....	10
What an Infinite Volume is .....	10
What a storage class is .....	11
What jobs are .....	11
What resource pools are .....	12
What rules and data policies are .....	12
Understanding Vserver associations .....	13
Database user capabilities .....	13
What availability health is .....	13
What mirror and backup vault protection relationships are .....	14
How protection relationships are created and protection jobs run .....	15
Tools for performing data restore operations .....	16
The virtual appliance backup and restore process overview .....	16
What events are .....	17
Event state definitions .....	17
<b>Common Unified Manager administrative workflows .....</b>	<b>19</b>
Configuring your environment after deployment .....	20
Changing the Unified Manager host name .....	21
Configuring Unified Manager to send alert notifications .....	25
Adding clusters .....	33
Changing the local user password .....	34
Monitoring and troubleshooting data availability .....	34
Resolving a flash card offline condition .....	35
Scanning for and resolving storage failover interconnect link down conditions .....	37
Resolving volume offline issues .....	40

Setting up and monitoring a Vserver with Infinite Volume without storage classes .....	45
Adding clusters .....	46
Editing the Infinite Volume threshold settings .....	47
Managing your Infinite Volume with storage classes and data policies .....	47
Editing the threshold settings of storage classes .....	49
Adding an alert .....	50
Creating rules .....	52
Exporting a data policy configuration .....	54
Resolving capacity issues .....	55
Performing suggested remedial actions for a full volume .....	56
Monitoring and troubleshooting protection relationships .....	57
Resolving a protection job failure .....	57
Resolving lag issues .....	61
Sending a support bundle to technical support .....	63
Accessing the maintenance console using Secure Shell .....	64
Generating a support bundle .....	64
Retrieving the support bundle using a Windows client .....	65
Retrieving the support bundle using a UNIX or Linux client .....	65
Sending a support bundle to technical support .....	67
Related tasks and reference information .....	67
Adding and reviewing notes about an event .....	67
Assigning events .....	68
Acknowledging and resolving events .....	68
Event details page .....	69
Description of event severity types .....	71
Description of event impact levels .....	72
Description of event impact areas .....	72
Volume details page .....	73
Vserver details page .....	83
Cluster details page .....	94
Aggregate details page .....	102
Job details page .....	107
Definitions of user roles in Unified Manager .....	108
Definitions of user types .....	108
Unified Manager roles and capabilities .....	109

<b>Using the maintenance console .....</b>	<b>111</b>
What the maintenance console does .....	111
What the maintenance user does .....	111
Diagnostic user capabilities .....	112
Accessing the maintenance console using Secure Shell .....	112
Accessing the maintenance console using the vSphere VM console .....	113
Maintenance console menu .....	113
Network Configuration menu .....	114
System Configuration menu .....	115
Support and Diagnostics menu .....	115
Adding additional network interfaces .....	116
<b>Troubleshooting Unified Manager issues .....</b>	<b>118</b>
VMware vSphere showing that VMware Tools are out-of-date .....	118
Remote User option does not display in the Add User dialog box .....	118
Alerts are not received by designated recipients .....	118
<b>Glossary .....</b>	<b>120</b>
<b>Copyright information .....</b>	<b>133</b>
<b>Trademark information .....</b>	<b>134</b>
<b>How to send your comments .....</b>	<b>135</b>
<b>Index .....</b>	<b>136</b>

# Unified Manager product and *Administration Guide* overview

---

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

## **Unified Manager web UI**

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

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

## **maintenance console**

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

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

# Concepts related to working with Unified Manager

---

Working with Unified Manager 6.0 involves understanding several concepts of which you should be aware.

You can also use the glossary to gain a better understanding of the terminology involved to describe Unified Manager concepts.

## Related concepts

*What a cluster is* on page 7

*What Vservers are* on page 8

*How volumes work* on page 9

*What a FlexVol volume is* on page 9

*Capabilities that FlexVol volumes provide* on page 10

*What an Infinite Volume is* on page 10

*What a storage class is* on page 11

*What jobs are* on page 11

*What resource pools are* on page 12

*What rules and data policies are* on page 12

*Understanding Vserver associations* on page 13

*Database user capabilities* on page 13

*What availability health is* on page 13

*What mirror and backup vault protection relationships are* on page 14

*How protection relationships are created and protection jobs run* on page 15

*Tools for performing data restore operations* on page 16

## Related references

*Glossary* on page 120

## 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.

A cluster can contain up to 24 nodes (unless the iSCSI or FC protocols are enabled, in which case the cluster can contain up to eight nodes). 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*.

### Related concepts

[Concepts related to working with Unified Manager](#) on page 7

## What Vservers are

A Vserver represents a single file-system namespace. It has separate network access and provides the same flexibility and control as a dedicated node. Each Vserver has its own user domain and security domain that can span multiple physical nodes.

A Vserver has a root volume that constitutes the top level of the namespace hierarchy; additional volumes are mounted to the root volume to extend the namespace. A Vserver is associated with one or more logical interfaces through which clients access the data on the storage system (or Vserver). Clients can access the Vserver from any node in the cluster through the logical interfaces that are associated with the Vserver.

**Note:** A namespace provides a context for determining the junctions that link together a collection of volumes. All the volumes associated with a Vserver are accessed from the Vserver's namespace.

### Related concepts

[Concepts related to working with Unified Manager](#) on page 7



## 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, 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 Vserver. The Vserver 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 Vserver it is associated with. The type of the volume (FlexVol volume or Infinite Volume) is determined by an immutable Vserver 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 Vserver.

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 a Vserver, then only those aggregates can be used to provide storage to the volumes associated with that Vserver. This impacts volume creation, and also copying and moving FlexVol volumes between aggregates.

For more information about Vservers, 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*.

### Related concepts

[Concepts related to working with Unified Manager](#) on page 7

## What a FlexVol volume is

A FlexVol volume is a data container associated with a Vserver with FlexVol volumes. It gets its storage from a single associated aggregate, which it might share with other FlexVol volumes or

Infinite Volumes. It can be used to contain files in a NAS environment, or LUNs in a SAN environment.

## Capabilities that FlexVol volumes provide

FlexVol volumes enable you to partition your data into individual manageable objects that can be configured to suit the needs of the users of that data.

A FlexVol volume enables you to take the following actions:

- Create a clone of the volume quickly and without having to duplicate the entire volume by using FlexClone technology.
- Reduce the space requirements of the volume by using deduplication and compression technologies.
- Create a Snapshot copy of the volume for data protection purposes.
- Limit the amount of space a user, group, or qtree can use in the volume by using quotas.
- Partition the volume by using qtrees.
- Create load-sharing mirrors to balance loads between nodes.
- Move the volume between aggregates and between storage systems.
- Make the volume available to client access using any file access protocol supported by Data ONTAP.
- Set up a volume to make more storage available when it becomes full.
- Create a volume that is bigger than the physical storage currently available to it by using thin provisioning.

## 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

*[Concepts related to working with Unified Manager](#) on page 7*

## 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 needs 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 performance storage class that uses only aggregates with SAS disks and the following volume settings: thin provisioning with compression and deduplication enabled.

### Related concepts

[Concepts related to working with Unified Manager](#) on page 7

## What jobs are

A job is a series of protection-related tasks initiated by a partner application that you can monitor using Unified Manager.

The following is a list of jobs you can monitor in Unified Manager:

- Restore
- Storage Service Cleanup
- Storage Service Conform
- Storage Service Destroy
- Storage Service Import
- Storage Service Modify
- Storage Service Subscribe
- Storage Service Unsubscribe
- Storage Service Update

### Related concepts

[Concepts related to working with Unified Manager](#) on page 7

## 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.

### Related concepts

*Concepts related to working with Unified Manager* on page 7

## What rules and data policies are

A *rule* determines the placement of files (data) in a Vserver 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 Vserver 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

*Concepts related to working with Unified Manager* on page 7

## Understanding Vserver associations

Virtual storage server (Vserver) associations are mappings from a source Vserver to a destination Vserver that are used by partner applications for resource selection and secondary volume provisioning.

You can associate Vservers in two ways:

- Associate any Vserver

You can create an association between any Vserver source to one or more destination Vservers. This means that all existing Vservers that currently require protection, as well as any Vservers that are created in the future, are associated with the specified destination Vservers. For example, you might want applications from several different sources at different locations backed up to one or more destination Vservers in one location.

- Associate a particular Vserver

You can create an association of a specific source Vserver with one or more destination Vservers. 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 Vserver source to a specific Vserver destination that is assigned to only that client.

### Related concepts

[Concepts related to working with Unified Manager](#) on page 7

## Database user capabilities

A database user can view data in the Unified Manager database. A database user does not have access to the Unified Manager web UI, maintenance console, and cannot execute API calls.

### Related concepts

[Concepts related to working with Unified Manager](#) on page 7

## What availability health is

*Availability health* is the reliability with which stored data can be accessed by authorized users. *Availability events* are events that indicate any hardware or software resource condition that impedes or blocks access to stored data by authorized users.

Unified Manager periodically monitors the hardware and software objects in your domain for conditions that adversely affect availability to your stored data.

Based on the monitored results, the Availability dashboard panel on the Dashboard page displays a graphic summary of your storage network's overall availability health and also displays the most recent or frequent events that might adversely affect availability of your stored data.

The event pages, inventory pages, and detail pages of Unified Manager web UI provide you with information to enable you to diagnose and identify the conditions that the availability events inform you of.

### Related concepts

*Concepts related to working with Unified Manager* on page 7

### Related tasks

*Resolving a flash card offline condition* on page 35

*Scanning for and resolving storage failover interconnect link down conditions* on page 37

*Resolving volume offline issues* on page 40

## What mirror and backup vault protection relationships are

Mirror protection relationships and backup vault protection relationships are protection configurations in which data stored in a source volume is protected by being replicated or backed up to a destination volume located in either the same storage cluster or a different storage cluster.

### Mirror protection (requires an active SnapMirror license)

In a mirror protection relationship, Snapshot copies of data in the source volume are replicated to a partner destination volume that is configured to be capable of taking over the data-serving functions of its partner source volume if that volume becomes unavailable. Mirror protection is enabled by activating the SnapMirror licenses on each cluster node.

### Backup vault protection (requires an active SnapVault license)

In a backup vault protection relationship, Snapshot copies of data in the source volume are backed up to a partner destination volume that is capable of providing storage-efficient and long-term retention of the backed up data. Backup vault protection is enabled by activating SnapVault licenses on each cluster node.

### Related concepts

*Concepts related to working with Unified Manager* on page 7

## How protection relationships are created and protection jobs run

Unified Manager enables you to discover, monitor, troubleshoot, and manage event resolution for existing mirror and backup vault protection relationships configured using various cluster management tools.

The following management tools enable a cluster or Vserver administrator to configure and run mirror or backup vault protection:

- **OnCommand System Manager**  
OnCommand System Manager 3.0 or later provides a user interface of windows and guided prompts to configure and execute mirror or backup vault protection.  
For more information, see the *OnCommand System Manager 3.0 Help*, which you can access from within OnCommand System Manager.
- **Clustered ONTAP SnapMirror Toolkit**  
The Clustered ONTAP SnapMirror Toolkit provides scripts to simplify pre-setup, setup, and initialization of mirror and vault protection relationships in a cluster environment.  
For more information on the toolkit or to obtain the toolkit itself, see the support download page, Clustered ONTAP SnapMirror Toolkit at <http://support.netapp.com/NOW/download/tools/smtk/>.
- **OnCommand Workflow Automation (WFA) 2.1 or later**  
OnCommand Workflow Automation provides a user interface with pre-configured workflows, which can be executed to configure mirror or backup vault protection.  
For more information, see the OnCommand Workflow Automation Help.  
For a PDF copy of the Workflow Automation help, see the product documentation page OnCommand Workflow Automation Product Library at <http://support.netapp.com/documentation/productlibrary/index.html?productID=61550>.
- **Data ONTAP CLI commands**  
The Data ONTAP CLI provides commands to configure and execute mirror or backup vault protection.  
For more information about creating protection relationships and implementing protection jobs through the Data ONTAP CLI, see the documentation about mirror and backup vault protection in the *Clustered Data ONTAP Data Protection Guide*.  
For a PDF copy of this guide, see the Data ONTAP 8 Product Library at <http://support.netapp.com/documentation/productlibrary/index.html?productID=30092>.

### Related concepts

[Concepts related to working with Unified Manager](#) on page 7

## Tools for performing data restore operations

Although Unified Manager enables you to monitor restore operations and recovery of data stored on managed clusters, it does not enable you to directly restore and recover data. However, you can use other cluster management tools to perform these operations.

The following management tools enable you to restore or recover data:

- OnCommand System Manager enables you to restore data using a web-based GUI interface. For more information, see the *OnCommand System Manager 3.0 Help*, which you can access from within OnCommand System Manager.
- Data ONTAP enables you to perform recovery and restore operations using the CLI. For more information, see the *Clustered Data ONTAP Data Protection Guide*. For a PDF copy of this guide, see the Data ONTAP 8 Product Library: <http://support.netapp.com/documentation/productlibrary/index.html?productID=30092>.

### Related concepts

*Concepts related to working with Unified Manager* on page 7

## The virtual appliance backup and restore process overview

The backup and restore model for Unified Manager is to capture and restore an image of the full virtual application.

The following tasks enable you to complete a backup of the virtual appliance:

1. Taking a VMware snapshot of the Unified Manager virtual appliance.
2. Making a NetApp Snapshot copy on the datastore to capture the VMware snapshot. If the datastore is not hosted on a Data ONTAP storage system, follow the storage vendor guidelines to create a backup of the VMware snapshot.
3. Replicating the NetApp Snapshot copy or snapshot equivalent to alternate storage.
4. Deleting the VMware snapshot.

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

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



## What events are

Events 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 they encompass. Impact areas include 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.

## 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 it is acknowledged.

**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 for 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, which in turn triggers the

Cluster Reachable event. Therefore, the Cluster Not Reachable event is no longer valid, and the event state changes to Obsolete.

# Common Unified Manager administrative workflows

---

Unified Manager is a storage monitoring and diagnostic tool. Some common administrative workflows associated with Unified Manager include selecting the storage clusters to monitor; diagnosing conditions that adversely affect data availability, capacity, and protection; configuration and management of infinite volumes; and, when necessary, bundling and sending diagnostic data to technical support.

Unified Manager is designed to enable a storage administrator to view a dashboard; assess the overall capacity, availability, and protection health of the managed storage clusters; and then quickly note, locate, diagnose, and assign for resolution any specific issues that might arise.

The most important cluster, Vserver, volume, infinite volume issue, or protection relationship issues that affect the storage capacity, data availability, or protection reliability of your managed storage are reflected in the Dashboard page system health graphs and posted events. When critical issues are signaled, the Dashboard page provides links to support appropriate troubleshooting workflows.

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

Common workflows relating to the following types of issues are described in this document:

- **Setting up the management environment after deployment**  
After storage clusters and their storage resources have been configured using the Data ONTAP CLI or System Manager, the storage administrator can further specify and configure those clusters for monitoring within Unified Manager.
- **Diagnosing and managing availability issues**  
If hardware failure or storage resource configuration issues cause the display of data availability events in the Dashboard page, the storage administrator can follow embedded links to display connectivity information on the affected storage resource, display troubleshooting advice, and assign issue resolution to other administrators.
- **Creating, configuring, monitoring, and protecting Infinite Volumes**  
After using the OnCommand workflow automation tool to create, configure, and define storage classes for an infinite volume, the storage administrator can use Unified Manager to monitor, set notification thresholds, and define data policy for that volume and its storage classes. Optionally, the storage administrator can use workflow automation and Unified Manager to set up data protection for the Infinite Volume.
- **Diagnosing and managing volume capacity issues**  
If volume storage capacity issues are reflected in the Dashboard page, a storage administrator can follow embedded links to display storage capacity current and historical trending information on the affected volume, display troubleshooting advice, and assign issue resolution to other administrators.
- **Monitoring and diagnosing protection relationship issues**

If protection reliability issues are displayed in the Dashboard page, a storage administrator can follow embedded links to display the current state of protection relationships, current and historical protection job success information about the affected relationships, and troubleshooting advice, and to assign issue resolution to other administrators.

- Sending a support bundle to technical support  
A storage administrator can retrieve and send a support bundle to technical support using the maintenance console. The administrator should send a support bundle when the issue in question requires more detailed diagnosis and troubleshooting than an AutoSupport message provides.

### Related concepts

[Monitoring and troubleshooting data availability](#) on page 34

[Monitoring and troubleshooting protection relationships](#) on page 57

### Related tasks

[Configuring your environment after deployment](#) on page 20

[Setting up and monitoring a Vserver with Infinite Volume without storage classes](#) on page 45

[Managing your Infinite Volume with storage classes and data policies](#) on page 47

[Resolving capacity issues](#) on page 55

[Sending a support bundle to technical support](#)

## 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 OnCommand Unified Manager Initial Setup.
- You must be logged in as the OnCommand Administrator to complete all tasks in this workflow.

### About this task

After you completed the OnCommand Unified Manager Initial Setup, you were presented with the option of adding 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 21

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 25

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](#) on page 33

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

### Related references

[Unified Manager roles and capabilities](#) on page 109

## 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 are 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 need to generate a new certificate if you change the host name. However, it is highly recommended that you do update the certificate, so that the host name in the certificate matches the actual host name.

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

### Steps

1. [Edit the network settings](#) on page 22

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 23

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 23

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 24

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

## 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### 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**

[Changing the Unified Manager host name](#) on page 21

## **Generating an HTTPS security certificate**

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

#### **Before you begin**

You must have the OnCommand Administrator role to perform 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 will take effect. This can be done from the System Configuration option in the NetApp maintenance console or from the VM console.

#### **After you finish**

After generating a new certificate, you can verify the new certificate information by viewing the HTTPS certificate.

#### **Related tasks**

[Changing the Unified Manager host name](#) on page 21

## **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 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 **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

[Changing the Unified Manager host name](#) on page 21

## 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.



## Related tasks

[Changing the Unified Manager host name](#) on page 21

## 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 signed in to the web UI as the OnCommand Administrator.

### 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 26

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 26

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 27

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

4. [Edit global threshold settings](#) on page 28

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 30

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 31

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

## Configuring notification settings

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

### Before you begin

The following information must be available:

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

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

### 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

[Configuring Unified Manager to send alert notifications](#) on page 25

## Enabling remote authentication

You can enable remote authentication (LDAP, Active Directory) to enable the management server to communicate with your authentication servers and to enable users of the authentication servers to use Unified Manager and manage the storage objects and data.

### Before you begin

You must have the role of OnCommand Administrator to enable remote authentication.

### About this task

If remote authentication is disabled, remote users or groups 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. Select **Enable Remote Authentication**.
4. Optional: Add authentication servers and test the authentication.
5. Click **Save and Close**.

### Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 25

[Adding authentication servers](#) on page 27

## 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

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.

### 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 tasks

[Configuring Unified Manager to send alert notifications](#) on page 25

## 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 29  
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 29  
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 30  
You can increase or decrease the warning or error lag time percentage so that events are generated at intervals that are more appropriate to your needs.

## Related tasks

*Configuring Unified Manager to send alert notifications* on page 25

## Configuring global aggregate threshold values

You can configure the global threshold values for all aggregates to track any threshold breach. Appropriate events are generated for threshold breaches and, based on these events, you can take preventive measures. 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### 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 > Aggregates**.
3. Configure the appropriate threshold values for capacity, growth, and Snapshot copies.
4. Click **Save and Close**.

## 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, based on these events, you can take preventive measures. 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### 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**.

## 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### 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**.

## 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 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.
4. Click **Add**.

#### Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 25

## 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 the notification settings such as 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 have the role of OnCommand Administrator to add an alert.

#### 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 “vol0” 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 **vol0** in the **Name contains** field to display the volumes whose name contains vol0.
  - b. Select <<**All Volumes whose name contains 'vol0'**>> 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.
  - a. Select **All Critical Events** from the Matching Events area and move it to the Selected Events area.
4. Click **Recipients** and enter **sample@domain.com** in the **Alert these users** field.
5. 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.

6. Click **Save**.

### Related concepts

[Event state definitions](#) on page 17

### Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 25

### Related references

[Description of event severity types](#) on page 71

[Description of event impact levels](#) on page 72

## 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.

### 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.  
The cluster management IP address must be the cluster-management LIF of the admin Vserver. If you use a node-management LIF, the operation fails.
  - User name and password to access the cluster
  - Type of protocol that can be configured on the cluster and the port number of the cluster
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### Steps

1. Click **Storage > Clusters**.
2. On the **Clusters** page, click **Add**.
3. In the **Add Cluster** dialog box, specify the values as required and then click **Add**.

### Result

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

## 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 OnCommand 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**.

## Monitoring and troubleshooting data availability

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

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

### Related concepts

*What availability health is* on page 13

### Related tasks

*Resolving a flash card offline condition* on page 35

*Scanning for and resolving storage failover interconnect link down conditions* on page 37

*Resolving volume offline issues* on page 40

## Resolving a flash card offline condition

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

### Before you begin

The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

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

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

### Steps

1. From the **Dashboard > Incidents and Risks > Availability Incidents** area, you click the hypertext link displayed for Flash Cards Offline.  
The Event details page for the availability incident displays.
2. On the **Event details** page, you can review the information displayed in the Cause field and perform one or more of the following tasks:
  - Assign the event to an administrator. [Assigning events](#) on page 68
  - Click the source of the event, in this case the cluster node on which the offline flash card is located, to get more information about that node. [Performing corrective action for a flash card offline](#) on page 35
  - Acknowledge the event. [Acknowledging and resolving events](#) on page 68

## Performing corrective action for a flash card offline

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

### Before you begin

The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

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

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

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

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

### Steps

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

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

In this example, the relevant information is in the Events summary table on the HA Details tab. The table specifies the flash card offline event, the time the event was generated, and, again, the cluster node from which this event originated.

2. Using the Data ONTAP CLI or System Manager, access the Event Manager System (EMS) logs for the affected cluster.

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

### After you finish

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

### Related references

[Event details page](#) on page 69

[Cluster details page](#) on page 94

[Unified Manager roles and capabilities](#) on page 109

## Scanning for and resolving storage failover interconnect link down conditions

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

### Before you begin


The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

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

### Steps

1. To check for recent availability events related to storage failover issues, check the Availability Incidents section and the Availability Risks listings on the **Dashboard** page.
2. To check further for all availability events related to storage failover issues, perform the following steps:
  - a) Click the **Availability Incidents** link on the **Dashboard** page.  
The Events page displays all events on the monitored clusters.
  - b) On the **Events** page, select the options **Incident** and **Risk** in the Filter column.
  - c) Then, at the top of the **Events** page Names column, click  and enter **\*failover** in the text box to limit the event to display to storage failover-related events.  
All past events related to storage failover conditions are displayed.

### Example

In this scenario, the Unified Manager displays the event, Storage Failover Interconnect One or More Links Down in its Availability Incidents section.

3. If one or more events related to storage failover display either on the **Dashboard** page or on the **Events** page, perform the following steps:
  - a) Click the event title link to display event details for that event.

### Example

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

The Event details page for that event displays.

- b) On the **Event details** page, you can perform one or more of the following tasks:
  - Review the error message in the Cause field and evaluate the issue. [Performing corrective action for storage failover interconnect links down](#) on page 38
  - Assign the event to an administrator. [Assigning events](#) on page 68
  - Acknowledge the event. [Acknowledging and resolving events](#) on page 68

### Related references

[Event details page](#) on page 69

[Unified Manager roles and capabilities](#) on page 109

## Performing corrective action for storage failover interconnect links down

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

### Before you begin

The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

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

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

```
Severity: Warning
State: New
Impact Level: Risk
Impact Area: Availability
Source: aardvark
Source Type: Node
Acknowledged By:
Resolved By:
Assigned To:
Cause: At least one storage failover interconnected link
       between the nodes aardvark and bonobo is down.
       RDMA interconnect is up (Link0 up, Link1 down)
```

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

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

### Steps

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

### Example

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

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

### Example

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

### After you finish

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

### Related references

[Event details page](#) on page 69

[Cluster details page](#) on page 94

[Unified Manager roles and capabilities](#) on page 109

## Resolving volume offline issues

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

### Before you begin

The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

Volumes might be reported offline for several reasons:

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

You can use the Dashboard page and the Cluster, Server, and Volume details pages to confirm or eliminate one or more of these possibilities.

### Steps

1. From the **Dashboard > Incidents and Risks > Availability Incidents** area, you click the Volume Offline event title text.

The Event details page for the availability incident displays.

2. On that page, check the notes for any indication that the Vserver administrator has taken the volume in question offline.
3. On the **Event details** page, you can review the information for one or more of the following tasks:
  - Review the information displayed in the Cause field for possible diagnostic guidance. In this example, the information in the Cause field informs you only that the volume is offline.



- Check the Notes and Updates area for any indication that the Vserver administrator has deliberately taken the volume in question offline.
- Click the source of the event, in this case the volume that is reported offline, to get more information about that volume. [Performing corrective action for volume offline conditions](#) on page 41
- Assign the event to an administrator. [Assigning events](#) on page 68
- Acknowledge the event or, if appropriate, mark it as resolved. [Acknowledging and resolving events](#) on page 68

## Performing diagnostic actions for volume offline conditions

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

### Before you begin

The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

Assuming that the volume that is reported offline has not been taken deliberately offline, that volume might be offline for several reasons.

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

### Choices

- Click **Volume details** page links to determine if the volume is offline because its host cluster node is down and storage failover to its HA pair partner has failed also. See [Determining if a volume offline condition is caused by a down cluster node](#) on page 42.
- Click **Volume details** page links to determine if the volume is offline and its host Vserver is stopped because the cluster node hosting the Vserver's root volume is down. See [Determining if a volume is offline and Vserver is stopped because a cluster node is down](#) on page 43.
- Click **Volume details** page links to determine if the volume is offline because of broken disks in its host aggregate. See [Determining if a volume is offline because of broken disks in an aggregate](#) on page 44.

### Related references

[Unified Manager roles and capabilities](#) on page 109

[Volume details page](#) on page 73

[Vserver details page](#) on page 83

[Cluster details page](#) on page 94

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

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

### Before you begin

The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

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


### Steps

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

The Vserver details page displays information about the offline volume's hosting Vserver.

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

The Volumes page displays a table of information about all the volumes hosted by the Vserver.

3. On the **Volumes** page State column header, click the filter symbol  and then select the option **Offline**.

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

4. On the **Volumes** page, click the grid symbol  and then select the option **Cluster Node**.

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

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

5. On the **Volumes** page, locate the listing for the offline volume and, in its Cluster Node column, click name of its hosting cluster node.

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

**After you finish**

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

**Determining if a volume is offline and its Vserver is stopped because a cluster node is down**

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because its host Vserver is stopped due to the cluster node hosting the Vserver's root volume being down.

**Before you begin**

The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.


**About this task**

To determine if the volume offline condition is caused its host Vserver being stopped because the cluster node hosting the Vserver's root volume is down, perform the following actions:

**Steps**

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

The Vserver details page displays the “running” or the “stopped” status of the hosting Vserver. If the Vserver status is running, then the volume offline condition is not caused by the cluster node hosting the Vserver's root volume being down.

3. If the Vserver status is stopped, then click **View Vservers** to further identify the cause of the hosting Vserver being stopped.
4. On the **Vservers** page Vserver column header, click the filter symbol  and then type the name of the stopped Vserver.

The information for that Vserver is shown in a table.

5. On the **Vservers** page, click the grid symbol  and then select the option **Root Volume**.

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

6. In the Root Volume column, click the name of the root volume to display the **Vserver details** page for that volume.  
If the status of the Vserver root volume is (Online), then the original volume offline condition is not caused because the cluster node hosting the Vserver's root volume is down.
7. If the status of the Vserver root volume is (Offline), then locate and click the hypertext link displayed under Aggregate in the **Related Devices** pane of the Vserver root volume's **Volume details** page
8. Locate and click the hypertext link displayed under Node in the **Related Devices** pane of the Aggregate's **Aggregate details** page  
The Nodes tab on the Cluster details page displays the state of the HA pair of nodes to which the Vserver root volume's hosting cluster node belongs. The state of the cluster node is indicated in the display.

### After you finish

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

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

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

### Before you begin

The following requirements must be met:

- Unified Manager must be deployed.
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

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

### Steps

1. Locate and click the hypertext link displayed under Aggregate in the **Related Devices** pane of the **Volume details** page.  
The Aggregate details page displays the online or the offline status of the hosting aggregate. If the aggregate status is online, then RAID disk problems are not the cause of the volume being offline.
2. If the aggregate status is offline, then click **Disk Information** and look for broken disk events in the **Events** list on the **Disk Information** tab.

3. To further identify the broken disks, click the hypertext link displayed under Cluster in the **Related Devices** pane.

The Cluster details page displays.

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

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

### After you finish

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

## Setting up and monitoring a Vserver with Infinite Volume without storage classes

You should use OnCommand Workflow Automation (WFA) and Unified Manager to set up and monitor a Vserver with Infinite Volume. You should create the Vserver with Infinite Volume using WFA and then monitor the Infinite Volume using Unified Manager. Optionally, you can configure data protection for your Infinite Volume.

### Before you begin

The following requirements must be met:

- Unified Manager must be deployed.
- WFA must be installed and the data sources must be configured.

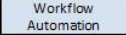
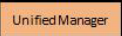
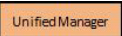

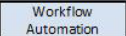
**Important:** WFA must be installed on a separate server.

- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.
- You must have created the required number of aggregates by customizing the appropriate predefined workflow in WFA.
- You must have configured Unified Manager server as a data source in WFA, and then you must have verified that the data is cached successfully.

### About this task

- You can monitor only data Vservers using Unified Manager.
- While performing this task, you are required to switch between two applications: OnCommand Workflow Automation (WFA) and OnCommand Unified Manager.
- The task provides high-level steps. For details about performing the WFA tasks, see the *OnCommand Workflow Automation* documentation.

## Steps

1.  Create a Vserver with Infinite Volume, and then create the Infinite Volume by using the appropriate workflow.  
You can enable storage efficiency technologies, such as deduplication and compression, while creating the Infinite Volume.
2.  Add the cluster containing the Vserver with Infinite Volume to the Unified Manager database.  
You can add the cluster by providing the IP address or the FQDN of the cluster.
3.  Based on your organization's requirements, modify the thresholds for the Infinite Volume on the Vserver.  
You should use the default Infinite Volume threshold settings.
4.  Configure notification alerts and traps to address any availability and capacity issues related to the Infinite Volume.
5. Optional:  Create a disaster recovery (DR) Vserver with Infinite Volume, and then configure data protection (DP) by performing the following steps:
  - a) Create a data protection (DP) Infinite Volume by using the appropriate workflow.
  - b) Set up a DP mirror relationship between the source and destination by using the appropriate workflow.

## Related tasks

[Adding clusters](#) on page 33

[Editing the Infinite Volume threshold settings](#) on page 47

[Adding an alert](#) on page 31

## Related references

[Unified Manager roles and capabilities](#) on page 109

## 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.

### 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.

The cluster management IP address must be the cluster-management LIF of the admin Vserver. If you use a node-management LIF, the operation fails.

- User name and password to access the cluster
- Type of protocol that can be configured on the cluster and the port number of the cluster
- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### Steps

1. Click **Storage > Clusters**.
2. On the **Clusters** page, click **Add**.
3. In the **Add Cluster** dialog box, specify the values as required and then click **Add**.

### Result

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

## 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### Steps

1. Click **Storage > Vservers**.
2. In the **Vservers** page, select the required Vserver with Infinite Volume.
3. In the **Vserver details** page, click **Actions > Edit Thresholds**.
4. In the **Edit Vserver with Infinite Volume Thresholds** dialog box, modify the thresholds as required.
5. Click **Save and Close**.

## Managing your Infinite Volume with storage classes and data policies

You can effectively manage your Infinite Volume by creating the Infinite Volume with the required number of storage classes, configuring thresholds for each storage class, creating rules and a data

policy to determine the placement of data written to the Infinite Volume, configuring data protection, and optionally configuring notification alerts.

**Before you begin**

The following requirements must be met:

- Unified Manager must be deployed.
- OnCommand Workflow Automation (WFA) must be installed.

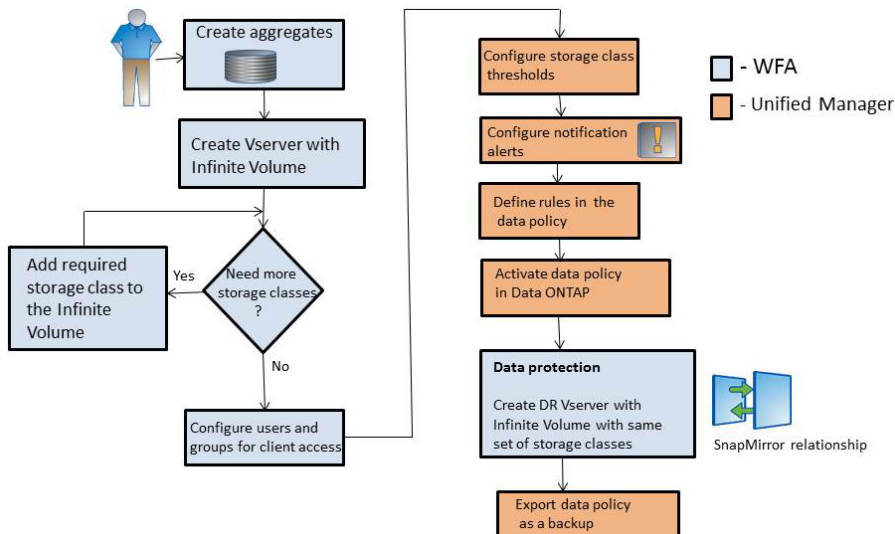
**Note:** Unified Manager and WFA are separate installations.

- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.
- You must have created the required number of aggregates by customizing the appropriate predefined workflow in WFA.
- You must have created the required number of storage classes by customizing the appropriate predefined workflow in WFA.
- You must have configured Unified Manager server as a data source in WFA, and then you must have verified that the data is cached successfully.

**About this task**

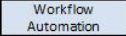
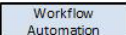
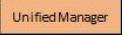
While performing this task, you are required to switch between two applications: OnCommand Workflow Automation (WFA) and OnCommand Unified Manager.

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





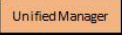
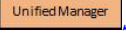
## Steps

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

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

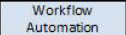
4.  *Based on your organization's requirements, modify the thresholds for each storage class* on page 49

You should use the default storage class threshold settings to effectively monitor storage class space.

5.  *Configure notification alerts and traps to address any availability and capacity issues related to the Infinite Volume* on page 31
6.  *Set up rules in the data policy, and then activate all the changes made to the data policy* on page 52

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

**Note:** Rules in a data policy affect only new data written to the Infinite Volume and do not affect existing data in the Infinite Volume.

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

## Editing the threshold settings of storage classes

When you need to address any issues related to storage space in your storage classes, you can edit the threshold settings of the storage class capacity based on your organization's requirements. When the threshold is crossed, events are generated, and you receive notifications if you have configured alerts for such events.

### Before you begin

You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### Steps

1. Click **Storage > Vservers**.
2. In the **Vservers** page, select the required Vserver with Infinite Volume.
3. In the **Vserver details** page, click **Actions > Edit Thresholds**.
4. In the **Edit Storage Class Thresholds** dialog box, modify the thresholds as required.
5. Click **Save and Close**.

## Adding 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 the notification settings such as 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 have the role of OnCommand Administrator to add an alert.

### 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 “vol0” 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 **vol0** in the **Name contains** field to display the volumes whose name contains vol0.
  - b. Select <<**All Volumes whose name contains 'vol0'**>> 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.
  - a. Select **All Critical Events** from the Matching Events area and move it to the Selected Events area.
4. Click **Recipients** and enter **sample@domain.com** in the **Alert these users** field.
5. 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.

6. Click **Save**.

### Related concepts

[Event state definitions](#) on page 17

### Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 25

### Related references

[Description of event severity types](#) on page 71

[Description of event impact levels](#) on page 72

## 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.
- The cluster containing the Vserver with Infinite Volume with storage classes must be added to the Unified Manager database.

### Choices

- [Creating rules using templates](#) on page 52
- [Creating custom rules](#) on page 53

## 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 Vserver with Infinite Volume. You can create rules based on file types, directory paths, or owners.

### Before you begin

- You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

- The cluster containing the Vserver 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 a Vserver with Infinite Volume.

### Steps

1. Click **Storage > Vservers**.
2. In the **Vservers** page, select the appropriate Vserver.
3. Click the **Data Policy** tab.

The list of rules in the data policy for the selected Vserver 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 Vserver.

### 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 Vserver 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.
- The cluster containing the Vserver 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 a Vserver with Infinite Volume.

**Steps**

1. Click **Storage > Vservers**.
2. In the **Vservers** page, select the appropriate Vserver.
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 Vserver.

**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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

**About this task**

The Data Policy tab, which is used while performing this task, is visible only for a Vserver with Infinite Volume.

**Steps**

1. Click **Storage > Vservers**.
2. In the **Vservers** page, select the appropriate Vserver.
3. Click **Data Policy**.

The list of rules in the data policy for the selected Vserver 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.

## Resolving capacity issues

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

### Before you begin

You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance. A user with any role can perform all of the tasks in this workflow that use Unified Manager.

### About this task

On the Dashboard page, you look at the Unresolved Incidents and Risks area and see a Volume Space Full error event in the Capacity pane under Vserver Volume Capacity at Risk.

### Steps

1. In the **Unresolved Incidents and Risks** area of the **Dashboard** page, click the name of the Volume Space Full error event in the **Capacity** pane.  
The Event details page for the error displays.
2. From the **Event details** page, you can perform one or more of the following tasks:
  - Review the error message in the Cause field and click on the suggestions under Suggested Remedial Actions to review descriptions of possible remediations. *Performing suggested remedial actions for a full volume* on page 56
  - Click the object name, in this case a volume, in the Source field to get details about the object. *Volume details page* on page 73
  - Look for notes that might have been added about this event. *Adding and reviewing notes associated with an event* on page 67
  - Add a note to the event. *Adding and reviewing notes associated with an event* on page 67
  - Assign the event to another user. *Assigning events* on page 68
  - Acknowledge the event. *Acknowledging and resolving events* on page 68
  - Mark the event as resolved. *Acknowledging and resolving events* on page 68

[Event details page](#) on page 69

## Performing suggested remedial actions for a full volume

After reviewing the suggested remedial actions on the Event details page for a Volume Space Full error event, you decide to perform one of the suggested actions.

### Before you begin

You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance. A user with any role can perform all of the tasks in this workflow that use Unified Manager.

### About this task

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

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

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

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

### Steps

1. From the **Event details** page, you click the volume name in the Source field to view details about the affected volume.
2. On the **Volume details** page, you click **Configuration** and see that deduplication and compression are already enabled on the volume.  
You decide to resize the volume.
3. In the **Related Devices** pane, you click the name of the hosting aggregate to see if the aggregate can accommodate a larger volume.
4. On the **Aggregate details** page, you see that the aggregate hosting the full volume does have enough uncommitted capacity, so you use OnCommand System Manager to resize the volume, giving it more capacity.

### Related references

[Event details page](#) on page 69

[Volume details page](#) on page 73

[Aggregate details page](#) on page 102



## Monitoring and troubleshooting protection relationships

Unified Manager enables you to monitor and troubleshoot mirror protection and backup vault protection of data stored on managed clusters.

### Related concepts

[How protection relationships are created and protection jobs run](#) on page 15

[Tools for performing data restore operations](#) on page 16

### Related tasks

[Resolving a protection job failure](#) on page 57

[Resolving lag issues](#) on page 61

## Resolving a protection job failure

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

### Before you begin

Because some tasks in this workflow require that you log in using the OnCommand Administrator role, you must be familiar with the roles required to use various functionality, as described in [Unified Manager roles and capabilities](#) on page 109.

### About this task

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

### Steps

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

**Tip:** The linked text for the event is written in the form `object_name:/object_name - Error Name`, such as `cluster2_src_vserver:/cluster2_src_vol2 - Protection Job Failed`.

The Event details page for the failed protection job displays.

2. Review the error message in the Cause field of the **Summary** area to determine the problem and evaluate potential corrective actions.

See [Identifying the problem and performing corrective actions for a failed protection job](#) on page 58.

### Related references

[Unified Manager roles and capabilities](#) on page 109

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

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

### Before you begin

You must be logged in as the OnCommand Administrator role to perform all tasks in this workflow.

### About this task

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

```
Protection Job Failed. Reason: (Transfer operation for
relationship 'cluster2_src_vserver:cluster2_src_vol2-
>cluster3_dst_vserver:
managed_svc2_vol3' ended unsuccessfully. Last error reported by
Data ONTAP: Failed to create Snapshot copy 0426cluster2_src_vol2snap
on volume cluster2_src_vserver:cluster2_src_vol2. (CSM: An operation
failed due to an ONC RPC failure..))
Job Details
```

This message provides the following information:

- A backup or mirror job did not complete successfully.  
The job involved a protection relationship between the source volume `cluster2_src_vol2` on the virtual server `cluster2_src_vserver` and the destination volume `managed_svc2_vol3` on the virtual server named `cluster3_dst_vserver`.
- A Snapshot copy job failed for `0426cluster2_src_vol2snap` on the source volume `cluster2_src_vserver:/cluster2_src_vol2`.

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

### Steps

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

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


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

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

The Volume details page displays for `cluster2_src_vol2`, showing the content of the Protection tab.

4. Looking at the protection topology graph, you see an error icon associated with the first volume in the topology, which is the source volume for the SnapMirror relationship.

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



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

Capacity information about volume `cluster2_src_vol2` displays.
7. In the **Capacity** pane, you see that there is an error icon in the bar graph, again indicating that the volume capacity has surpassed the threshold level set for the volume.

8. Below the capacity graph, you see that the autogrow option has been disabled and that a volume space guarantee has been set.

You could decide to enable autogrow, but for the purposes of this scenario, you decide to investigate further before making a decision about how to resolve the capacity problem.

9. You scroll down to the **Events** list and see that Protection Job Failed, Volume Days Until Full, and Volume Space Full events were generated.
10. In the **Events** list, you click the **Volume Space Full** event to get more information, having decided that this event seems most relevant to your capacity issue.

The Event details page displays the Volume Space Full event for the source volume.

11. In the **Summary** area, you read the Cause field for the event: The full threshold set at 90% is breached. 45.38 MB (95.54%) of 47.50 MB is used.
12. Below the **Summary** area, you see Suggested Corrective Actions.

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

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

(Each of the suggested tasks in this example must be performed either from the System Manager web UI or from the Data ONTAP CLI console.)

- Enable autogrow on this volume.
- Resize the volume.
- Enable and run deduplication on this volume.
- Enable and run compression on this volume.

13. You decide to enable autogrow on the volume, but to do so, you must determine the available free space on the parent aggregate and the current volume growth rate:

- a) Look at the the parent aggregate, `cluster2_src_aggr1`, in the **Related Devices** pane.

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

You determine that the aggregate has sufficient space to enable volume autogrow.

- b) At the top of the page, look at the icon indicating a critical incident and review the text below the icon.

You determine that "Days to Full: Less than a day | Daily Growth Rate: 5.4%".

14. Go to System Manager or access the Data ONTAP CLI to enable the volume autogrow option.

**Tip:** Make note of the names of the volume and aggregate so you have them available when enabling autogrow.

15. After resolving the capacity issue, return to the Unified Manager **Event details** page and mark the event as resolved.

## Related tasks

[Adding and reviewing notes about an event](#) on page 67

[Assigning events](#) on page 68

[Acknowledging and resolving events](#) on page 68

## Related references

[Job details page](#) on page 107

## Resolving lag issues

This workflow provides an example of how you might resolve a lag issue. In this scenario, you are an administrator or operator accessing the Unified Manager Dashboard page to see if there are any problems with your protection relationships and, if one exists, to find a solution.

### Before you begin

You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### About this task

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

### Steps

1. In the **Protection** pane on the **Dashboard** page, locate the SnapMirror relationship lag error and click it.


The Event details page for the lag error event is displayed.

2. From the **Event details** page you can perform one or more of the following tasks:
  - Review the error message in the Cause field of the Summary area to determine if there is any suggested corrective action.
  - Click the object name, in this case a volume, in the Source field of the Summary area to get details about the volume.
  - Look for notes that might have been added about this event.
  - Add a note to the event.
  - Assign the event to a specific user.
  - Acknowledge or resolve the event.
3. In this scenario, you click the object name (in this case, a volume) in the Source field of the **Summary** area to get details about the volume.

The Protection tab of the Volume details page is displayed.

4. In the **Protection** tab, you look at the topology diagram.

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



5. Click each of the volumes in the SnapMirror cascade.

As you select each volume, the protection information in the Summary, Topology, History, Events, Related Devices, and Related Alerts areas changes to display details relevant to the selected volume.

6. You look at the **Summary** area and position your cursor over the information icon in the **Update Schedule** field for each volume.

In this scenario, you note that the SnapMirror policy is DPDefault, and the SnapMirror schedule updates hourly at 5 minutes after the hour. You realize that all the volumes in the relationship are attempting to complete a SnapMirror transfer at the same time.

7. To resolve the lag issue, you modify the schedules for two of the cascaded volumes so that each destination begins a SnapMirror transfer after its source has completed a transfer.

You can use OnCommand System Manager to change the SnapMirror schedule for your volumes. For more information, see the *OnCommand System Manager Help*, which you can access from within OnCommand System Manager.

### Related tasks

- [Adding and reviewing notes about an event](#) on page 67
- [Assigning events](#) on page 68
- [Acknowledging and resolving events](#) on page 68

### Related references

- [Event details page](#) on page 69

*Unified Manager roles and capabilities* on page 109

## Sending a support bundle to technical support

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

### Before you begin

You must be the maintenance user to complete this workflow.

### About this task

For more information about the maintenance console and support bundles, see *Using the maintenance console* on page 111.

Unified Manager stores two generated support bundles at one time.

### Steps

1. *Accessing the maintenance console using Secure Shell* on page 64  
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.
2. *Generate a support bundle* on page 64  
You can generate a support bundle using the maintenance console. After you generate the support bundle, you need to retrieve it using either a Windows, Unix, or Linux client.
3. *Retrieve the support bundle using a Windows client* on page 65  
You can use a retrieval tool such as Filezilla or WinSCP to retrieve the support bundle. Alternatively, if you use a Unix or Linux client, you can retrieve the support bundle using the CLI.
4. *Retrieve the support bundle using a Unix or Linux client* on page 65  
If you use a Unix or Linux client, you can retrieve the support bundle using CLI. After retrieving the support bundle, you can upload it to the technical support website.
5. *Send the support bundle to technical support* on page 67  
You can upload the support bundle to technical support to receive additional troubleshooting help.

## 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 have the maintenance user role.

### 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.

### Related tasks

[Using the maintenance console](#) on page 111

## Generating a support bundle

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

### Before you begin

You must have accessed the maintenance console as the maintenance user.

### About this task

Unified Manager stores two generated support bundles at one time.

### Steps

1. From **Main Menu**, select **Support/Diagnostics menu**.
2. Select **Generate Support Bundle**.

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

### After you finish

After generating the support bundle, you can retrieve it using an SFTP client or by using UNIX or Linux CLI commands.



**Related concepts**

*Diagnostic user capabilities* on page 112

**Related references**

*Unified Manager roles and capabilities* on page 109

**Retrieving the support bundle using a Windows client**

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

**Before you begin**

You must be the maintenance user to perform this task.

You must use a tool that supports SCP or SFTP.

**Steps**

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

**Related information**

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

*WinSCP* - <http://winscp.net>

**Retrieving the support bundle using a UNIX or Linux client**

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

**Before you begin**

You must be the maintenance user to perform this task.

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

**Steps**

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

---

**If you are using... Then use the following command...**

---

SCP	<code>scp &lt;maintenance-user&gt;@&lt;vApp-name-or-ip&gt;:/support/ support_bundle_file_name.7z &lt;destination-directory&gt;</code>
-----	---

---

SFTP	<code>sftp &lt;maintenance-user&gt;@&lt;vApp-name-or-ip&gt;:/support/ support_bundle_file_name.7z &lt;destination-directory&gt;</code>
------	--

---

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

4. Enter the maintenance user password.

**Examples**

The following example uses SCP to retrieve the support bundle:

```
$ scp admin@10.228.212.69:/support/  
support_bundle_20130216_145359.7z .  
  
Password:  
maintenance_user_password  
support_bundle_20130216_145359.7z      100% 119MB 11.9MB/s   00:10
```

The following example uses SFTP to retrieve the support bundle:

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

## Sending a support bundle to technical support

When directed by technical support, you can send a support bundle using the direction provided in KB article 1010090. You should send a support bundle when the issue that you have requires more detailed diagnosis and troubleshooting than an AutoSupport message provides.

### Before you begin

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

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

### Steps

1. Go to the NetApp Support Site and log in.
2. Search for Knowledge Base article 1010090.
3. Follow the instructions on how to upload a file to technical support.

### Related information

[NetApp Support Site](#)

## Tasks and information related to several workflows

Some tasks and reference texts are common to many of the workflows in Unified Manager, including adding and reviewing notes about an event, assigning an event, and acknowledging and resolving events, and details about volumes, Vservers, aggregates, and so on.

### Adding and reviewing notes about an event

While addressing events, you can add information about how the issue is being addressed. This might enable another user who is assigned the event to address the event. The Notes and Updates area in the Event details page enables you to add such information. 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### 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**.

## 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 authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

### 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...
<b>Yourself</b>	Click <b>Assign To &gt; Me</b> .
<b>Another user</b>	<ol style="list-style-type: none"> <li>a. Click <b>Assign To &gt; Another user</b>.</li> <li>b. In the Assign Owner dialog box, enter the user name, or select a user from the drop-down list.</li> <li>c. Click <b>Assign</b>. An email notification is sent to the user.</li> </ol>

**Note:** If you do not enter a user name or select a user from the drop-down list, and click **Assign**, the event remains unassigned.

## Acknowledging and resolving events

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

### Before you begin

You must be authorized to perform all the steps of this task; your administrator can confirm your authorization in advance.

**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:

<b>If you want to...</b>	<b>Do this...</b>
Acknowledge and mark a single event as resolved	<ol style="list-style-type: none"> <li>a. Click the event name.</li> <li>b. From the Event details page, determine the cause of the event.</li> <li>c. Click <b>Acknowledge</b>.</li> <li>d. Take appropriate corrective action.</li> <li>e. Click <b>Mark As Resolved</b>.</li> </ol>
Acknowledge and mark multiple events as resolved	<ol style="list-style-type: none"> <li>a. Determine the cause of the events from the respective Event details page.</li> <li>b. Select the events.</li> <li>c. Click <b>Acknowledge</b>.</li> <li>d. Take appropriate corrective actions.</li> <li>e. Click <b>Mark As Resolved</b>.</li> </ol>

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

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

**Event details page**

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

- [Command buttons](#) on page 69
- [Summary area](#) on page 70
- [Notes and Updates area](#) on page 71
- [Suggested Corrective Actions area](#) on page 71

**Command buttons**

The command buttons enable you to perform the following tasks:

<b>Assign To</b>	<b>Me</b>	Assigns the event to you.
	<b>Another user</b>	<p>Opens the Assign Owner dialog box, which enables you to assign or reassign the event to other users.</p> <p>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.</p> <p><b>Note:</b> You can also unassign events by leaving the ownership field blank.</p>

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



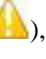

**Mark As Resolved** Enables you to change the event state to resolved.

**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:

<b>Severity</b>	Displays the severity of the event. The event severity types are Critical (  ) , Error (  ) , Warning (  ) , and Information (  ) .
<b>State</b>	Displays the event state: New, Acknowledged, Resolved, or Obsolete.
<b>Impact Level</b>	Displays whether the event is categorized as an incident, risk, or an informational event.
<b>Impact Area</b>	Displays whether the event is a capacity, availability, protection, or configuration related event.
<b>Obsoleted Cause</b>	Displays the reason the event is now obsolete.
<b>Source</b>	Displays the full name of the object along with the type of object with which the event is associated.
<b>Source Type</b>	Displays the object type (for example, Vserver, volume, or qtree) with which the event is associated.
<b>Acknowledged By</b>	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.

<b>Resolved By</b>	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.
<b>Assigned To</b>	Displays the name of the person who is assigned to work on the event.
<b>Cause</b>	Displays information about the cause of the event.
<b>Alert Settings</b>	<p>The following information about alerts is displayed:</p> <ul style="list-style-type: none"> <li>• If there are no alerts associated with the selected event, an <b>Add</b> link is displayed. You can open the Add Alert dialog box by clicking the link.</li> <li>• 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.</li> <li>• 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.</li> </ul> <p>Alerts that are disabled are not displayed.</p>

### 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.

### Suggested Corrective Actions area

Displays 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

[Performing diagnostic actions for volume offline conditions](#) on page 41

[Performing suggested remedial actions for a full volume](#) on page 56

## 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.

## 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 four impact areas (availability, capacity, configuration, 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 severity type of Information.
- Protection** Protection events notify you of incidents or risks involving SnapMirror lags or relationships, issues with destination capacity, problems with SnapVault relationships, or issues with protection jobs.



## 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, and events generated. You can also view information about the related objects and related alerts for that volume.

- [Command buttons](#) on page 73
- [Capacity tab](#) on page 73
- [Efficiency tab](#) on page 76
- [Configuration tab](#) on page 77
- [Protection tab](#) on page 78
- [History area](#) on page 82
- [Events list](#) on page 82
- [Related Devices pane](#) on page 82
- [Related Alerts pane](#) on page 83

### 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.

**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 data capacity graph and the Snapshot reserve graph, which display capacity details of the volume.

- **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 in the volume (free capacity).

**Snapshot Reserve** Displays the used and free Snapshot capacity of the volume.

**Note:** 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 autogrown. The default is 5% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.
- **Quota Committed Capacity**  
Displays the space reserved in the quotas.
- **Quota Overcommitted Capacity**  
Displays the amount of space that can be used before the system generates the Volume 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 due to lack of space in the aggregate.

- **Snapshot Count**  
Displays the number of Snapshot copies in the volume.

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

### **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.
  - **Vserver**  
Displays the name of the Vserver that contains the volume.
  - **Junction Path**  
Displays the status of the path, which can be active or inactive. The path in the Vserver 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 Vserver.
  - **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 due to lack of space in the aggregate.

- 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 lag duration since the last successful protection update.
- **Last Successful Update**  
Displays the date and time of the last successful protection update.
- **Managed**  
Displays either Yes or No to indicate whether or not the volume is managed by a storage service.
- **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.
- **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.
- **SnapMirror Policy**  
Displays the protection policy for the volume. DPDefault indicates the default SnapMirror protection policy, and XDPDefault indicates the default SnapVault policy. If XDPDefault is displayed, you can click the policy to view the rules associated with that policy.
- **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 blue 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.

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. 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 the graph you want to view from the drop-down list. You can also view details for a specified time period, such as one week, one month, or one year.

## Events list

The Events list displays details about new and acknowledged events:

<b>Severity</b>	Displays the severity of the event.
<b>Event</b>	Displays the event name.
<b>Triggered Time</b>	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 Vservers, aggregates, qtrees, LUNs, and Snapshot copies that are related to the volume:

<b>Vserver</b>	Displays the name, capacity, and the health status of the Vserver that contains the selected volume.
<b>Aggregate</b>	Displays the name, capacity, and the health status of the aggregate that contains the selected volume.
<b>Volumes in the Aggregate</b>	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.
<b>Qtrees</b>	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.
<b>NFS Exports</b>	Displays the number and status of the NFS exports associated with the volume.
<b>LUNs</b>	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.
<b>FlexClone Volumes</b>	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.
- FlexClone Volumes of the Parent Volume** Displays the number and capacity of all the clone volumes which belongs to the parent volume of the selected FlexClone volume. The number and capacity are 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.

### Related tasks

[Performing diagnostic actions for volume offline conditions](#) on page 41

[Performing suggested remedial actions for a full volume](#) on page 56

## Vserver details page

You can use the Vserver details page to view detailed information about the selected Vserver that is monitored by Unified Manager, such as its health, capacity, configuration, data policies, logical interfaces (LIFs), LUNs, and qtrees related to the Vserver. You can also view information about the related objects and related alerts for the Vserver.

**Note:** You can monitor only data Vservers.

- [Command buttons](#) on page 83
- [Health tab](#) on page 84
- [Capacity tab](#) on page 84
- [Configuration tab](#) on page 87
- [LIFs tab](#) on page 88
- [NFS Exports tab](#) on page 89
- [LUNs tab](#) on page 92
- [Qtrees tab](#) on page 90
- [Data Policy tab](#) on page 93
- [Related Devices pane](#) on page 94
- [Related Alerts pane](#) on page 94

### Command buttons

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

- Actions**
- Add Alert

- Enables you to add an alert to the selected Vserver.
- **Edit Thresholds**  
Enables you to edit Vserver thresholds.

**View Vservers** Enables you to navigate to the Vservers page.

## Health tab

The Health tab displays detailed information about data availability and data capacity issues of various objects such as volumes, aggregates, LIFs, LUNs, protocols, services, and NFS exports. Availability issues are related to the data serving capability of the Vserver objects. Capacity issues are related to the data storing capability of the Vserver objects:

**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 Vserver. For example, information is displayed about 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.

If the selected Vserver is a Vserver 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 Vserver. For example, information is displayed about aggregates that are likely to breach the set threshold values.

If the selected Vserver is a Vserver 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 Vserver.

The following information is displayed for a Vserver 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 Vserver.

- **Used**  
Displays the space used by data in the volumes that belong to the Vserver.
- **Guaranteed Available**  
Displays the guaranteed available space for data that is available for volumes in the Vserver.
- **Unguaranteed**  
Displays the available space remaining for data that is allocated for thinly provisioned volumes in the Vserver.

### **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). For example, red color indicates that the volume has breached the Space Full threshold, and yellow color indicates that the volume has breached the Space Nearly Full threshold.
- **Days to Full**  
Displays the estimated number of days remaining before the volume reaches full capacity.




The following information is displayed for a Vserver with Infinite volume:

<b>Capacity</b>	<p>Displays the following capacity-related details:</p> <ul style="list-style-type: none"> <li>• Percentage of used and free data capacity</li> <li>• Percentage of used and free Snapshot capacity</li> <li>• Used Displays the space used by data in the Vserver with Infinite Volume.</li> <li>• Warning Indicates that the space in the Vserver with Infinite Volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.</li> <li>• Error Indicates that the space in the Vserver with Infinite Volume is full. If this threshold is breached, the Space Full event is generated.</li> </ul>
<b>Other Details</b>	<ul style="list-style-type: none"> <li>• Total Capacity Displays the total capacity in the Vserver with Infinite Volume.</li> <li>• Data Capacity Displays used data capacity, available data capacity, and Snapshot overflow capacity details of the Vserver with Infinite Volume.</li> <li>• Snapshot Reserve Displays the used and free details of the Snapshot reserve.</li> <li>• System Capacity Displays the used system capacity and available system capacity in the Vserver with Infinite Volume.</li> <li>• Thresholds Displays the nearly full and full thresholds of the Vserver with Infinite Volume.</li> </ul>
<b>Storage Class Capacity Details</b>	<p>Displays information about the capacity usage in your storage classes. This information is displayed only if you have configured storage classes for your Vserver with Infinite Volume.</p>
<b>Vserver Storage Class Thresholds</b>	<p>Displays the following thresholds (in percentage) of your storage classes:</p> <ul style="list-style-type: none"> <li>• Nearly Full Threshold Specifies the percentage at which a storage class in a Vserver with Infinite Volume is considered to be nearly full.</li> <li>• Full Threshold Specifies the percentage at which the storage class in a Vserver with Infinite Volume is considered full.</li> <li>• Snapshot Usage Limit Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.</li> </ul>

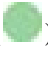

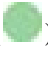

## Configuration tab

The Configuration tab displays configuration details about the selected Vserver, such as its cluster, root volume, the type of volumes it contains (Infinite Volume or FlexVol volumes), and the policies created on the Vserver:

### Overview

- Cluster  
Displays the name of the cluster to which the Vserver belongs.
- Allowed Volume Type  
Displays the type of volumes that can be created in the Vserver. The type can be InfiniteVol or FlexVol.
- Root Volume  
Displays the name of the root volume of the Vserver.
- Allowed Protocols  
Displays the type of protocols that can be configured on the Vserver. 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 Vserver. Also, indicates if the LIFs are up (  ) or down (  ).
- SAN  
Displays the number of SAN LIFs that are associated with the Vserver. 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 a Vserver with Infinite Volume only.
- Storage Classes  
Displays the storage classes associated with the selected Vserver with Infinite Volume. Storage classes are displayed for a Vserver with Infinite Volume only.

### Policies

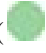


- Snapshots  
Displays the name of the Snapshot policy that is created on the Vserver.
- Export Policies  
Displays either the name of the export policy if a single policy is created or displays the number of export policies if multiple policies are created.
- Data Policy  
Displays whether a data policy is configured for the selected Vserver with Infinite Volume.

### Services

- Type

Displays the type of service that is configured on the Vserver. 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 Vserver:




### LIF

Displays the name of the LIF that is created on the selected Vserver.

### 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.

### Home Port

Displays the physical port to which the LIF was originally associated.

### Current Port



Displays the physical port to which the LIF is currently associated. If the LIF is migrated, the current port might be different from the home port.



<b>Failover Policy</b>	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.
<b>Routing Groups</b>	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.
<b>Failover Group</b>	Displays the name of the failover group.

## 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 in the following conditions: if the volume is not mounted or if the protocols associated with the export policy for the volume do not contain NFS.

<b>Status</b>	Displays the current status of the NFS export. The status can be Error (  ) or Normal (  )
<b>Junction Path</b>	Displays the path to which the volume is mounted.
<b>Junction Path Status</b>	Displays whether the path to access the mounted volume is active or inactive.
<b>Volume / Vserver</b>	Displays the name of the volume, if the volume being exported is a FlexVol volume. For Infinite Volumes, the name of the Vserver containing the Infinite Volume is displayed.
<b>Volume State</b>	Displays the state of the volume that is being exported. The state can be Offline, Online, or Restricted. <ul style="list-style-type: none"> <li>• Offline Read or write access to the volume is not allowed.</li> <li>• Online Read and write access to the volume is allowed.</li> <li>• Restricted Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.</li> </ul>
<b>Security Style</b>	Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed. <ul style="list-style-type: none"> <li>• UNIX (NFS clients) Files and directories in the volume have UNIX permissions.</li> <li>• Unified</li> </ul>





- 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.

<b>UNIX Permission</b>	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.
<b>Export Policy</b>	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.

### Qtrees tab

The Qtrees tab displays details about qtrees and their quotas:

**Note:** The Qtrees tab is not displayed for Vserver with Infinite Volume.

<b>Status</b>	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.

<b>Qtree</b>	Displays the name of the qtree.
<b>Volume</b>	Displays the name of the volume that contains the qtree.

<b>Quota Set</b>	Indicates whether a quota is enabled or disabled on the qtree.
<b>Disk Used (%)</b>	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.
<b>Disk Hard Limit</b>	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” in 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.
<b>Disk Soft Limit</b>	Displays the amount of disk space allocated for the qtree before a warning event is generated. The value is displayed as “Unlimited” in 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.
<b>Disk Threshold</b>	Displays the threshold value set on the disk space. The value is displayed as “Unlimited” in 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.
<b>Files Used (%)</b>	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.
<b>File Hard Limit</b>	Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” in 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.
<b>File Soft Limit</b>	Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” in 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.
<b>Filters pane</b>	Enables you to set filters to customize the display of information in the qtrees list. You can select filters in the Status column.  <b>Note:</b> The filters specified in the Filters pane override the filters specified for the columns in the qtrees list.

## Volume area

The Volume area provides more information about the selected qtree:

<b>Name</b>	Displays the name of the volume.
-------------	----------------------------------

<b>Used Capacity</b>	Displays, as a graph, information about the volume capacity usage.
<b>Quota Committed Capacity</b>	Displays the space reserved in the quotas.
<b>Space Guarantee</b>	Displays the flexible 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: <ul style="list-style-type: none"> <li><b>None</b> No space guarantee is configured for the volume.</li> <li><b>File</b> Full size of sparsely written files, for example LUNs, are guaranteed.</li> <li><b>Volume</b> Full size of the volume is guaranteed.</li> <li><b>Partial</b> 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. <ul style="list-style-type: none"> <li><b>Note:</b> The space guarantee is Partial when the volume is of type Data-Cache.</li> </ul> </li> </ul>

### Related Devices area





The Related Devices area enables you to view and navigate to the LUNs that are related to the qtree:

**LUNs** Displays the number of all the LUNs in the selected qtree. The health status of the LUNs is also displayed, based on the highest severity level.

### LUNs tab

Displays details about the LUNs that belong to the selected Vserver:

**Note:** The LUNs tab is not displayed for Vserver with Infinite Volume.

<b>Status</b>	Displays the current status of the LUN. The status can be Critical (  ) , Error (  ) , Warning (  ) , or Normal (  ) .
<b>LUN</b>	Displays the name of the LUN.
<b>State</b>	Displays whether the LUN is online or offline.
<b>Container Path</b>	Displays the name of the file system (volume or qtree) that contains the LUN.
<b>Total Size</b>	Displays the total data capacity of the LUN.

<b>Thin Provisioned</b>	Displays whether thin provisioning is enabled for the LUN.
<b>Serial Number</b>	Displays the serial number of the LUN.
<b>Initiator Groups area</b>	<p>The Initiator Groups area provides more information about the selected LUN:</p> <ul style="list-style-type: none"> <li>• <b>Initiator Group</b> Displays the initiator groups that the LUN is mapped to.</li> <li>• <b>OS Type</b> Displays the type of host operating system used by all of the initiators in the group.</li> <li>• <b>Initiator Group Type</b> Displays the type of the initiators in the group. Possible values are: iSCSI, FC/FCoE, and Mixed.</li> </ul>

### 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 Vservers 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 Vserver 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 pane

The Related Devices pane enables you to view the cluster, aggregates, and volumes that are related to the Vserver:

- Cluster** Displays the name and health status of the cluster to which the Vserver belongs.
- Aggregates** Displays the number of aggregates that belong to the selected Vserver. The health status of the aggregates is also displayed, based on the highest severity level. For example, if a Vserver 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 that belong to the selected Vserver. 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 Vserver. You can also add an alert by clicking the **Add Alert** link or edit an existing alert by clicking the alert name.

### 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 95
- *Health tab* on page 95
- *Capacity tab* on page 96
- *Configuration tab* on page 97
- *LIFs tab* on page 98
- *Nodes tab* on page 99
- *Disks tab* on page 100
- *Related Devices pane* on page 102
- *Related Alerts pane* on page 102

## Command buttons

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

- |                      |   |
|----------------------|---|
| <b>Actions</b>       | <ul style="list-style-type: none"> <li>• Add Alert<br/>Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.</li> <li>• Edit Cluster<br/>Opens the Edit Cluster dialog box, which enables you to edit the settings of the selected cluster.</li> </ul> |
| <b>View Clusters</b> | 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, Vservers, 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.

- 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.
- **Spares**  
Displays the storable capacity (in MB, GB, and so on) available for storage in all the spare disks.
- **Provisioned**  
Displays the capacity (in MB, GB, and so on) that is provisioned for all the underlying volumes.

**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). For example, red color indicates that the aggregate has breached the Space Full threshold, and yellow color indicates that the aggregate has breached the Space Nearly Full threshold. If the aggregate is overcommitted, the committed capacity is also displayed.
- **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.
  - **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.
- Vservers**
- **Availability**  
Displays the number of Vservers 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. This field is blank if the required protocol licenses are not enabled or when the node locked licenses are enabled.

## 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.

<b>Administrative Status</b>	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.
<b>IP Address</b>	Displays the IP address of the LIF.
<b>Role</b>	Displays the role of the LIF. Possible roles are Cluster-Management LIFs, Node-Management LIFs, Cluster LIFs, and Intercluster LIFs.
<b>Home Port</b>	Displays the physical port to which the LIF was originally associated.
<b>Current Port</b>	Displays the physical port to which the LIF is currently associated. After LIF migration, the current port might be different from the home port.
<b>Failover Policy</b>	Displays the failover policy that is configured for the LIF.
<b>Routing Groups</b>	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.
<b>Failover Group</b>	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, and service processor. 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 associated with the HA pair. You can view environmental details of the disk shelves, such as power supplies, fans, temperature sensors, current sensors, and voltage sensors. These environmental details are represented in terms of icons that are color-coded:

- Green: The environmental components are in working condition.
- Grey: There is no data available for the environmental components.
- Red: Some of the environmental components are down.

You can view other disk shelves related information, such as the state of the disk shelf, the model number, the unique ID, and the firmware revision number.

You can also view events related to the disk shelves and the environmental components, and the time when the events were triggered.

**Ports** Displays information about the associated ports:

<b>Port</b>	Displays the name of the port. For example, the port names can be e0M, e0a, and e0b.
<b>State</b>	Displays the current state of the port. The possible states are (  ), Down (  ), or Unknown (  ). By default, this column is sorted to display the states in the following order: Down, Up, and Unknown.
<b>Owner Node</b>	Displays the name of the node to which the port belongs.
<b>Type</b>	Displays the protocol used for the port. The possible types are Ethernet, Fibre Channel, and FCoE.
<b>Role</b>	Displays the role of the port. The possible roles are Cluster, Data, Intercluster, Node-Management, and Undefined.
<b>LIF</b>	Displays the number of LIFs associated with the port.  By clicking the number, you can view LIF details such as the LIF name, status, network address of the LIF, protocol, current port, and the Vserver associated with the LIF.

## 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:

<b>Disk Pool Summary</b>	Displays the number of disks categorised by effective types (FCAL, SAS, SATA, and SSD) and the state of the disks. You can also view other details of the disks such as the number of broken disks, spare disks, and unassigned disks.
<b>Disk</b>	Displays the name of the disk.
<b>RAID Groups</b>	Displays the name of the RAID group.
<b>Owner Node</b>	Displays the name of the node to which the disk belongs. If the disk is unassigned, no value is displayed in this column.
<b>State</b>	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.
<b>Position</b>	Displays the position of the disk based on its container type—for example, Copy, Data, or Parity. By default, this column is hidden.
<b>Impacted Aggregate</b>	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.
<b>Storable Capacity</b>	Displays the disk capacity that is available for use.
<b>Raw Capacity</b>	Displays the capacity of the raw, unformatted disk before right-sizing and RAID configuration. By default, this column is hidden.
<b>Type</b>	Displays the types of disks—for example, ATA, SATA, or FCAL.
<b>Effective Type</b>	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.
<b>Firmware</b>	Displays the firmware version of the disk.
<b>RPM</b>	Displays the revolutions per minute (RPM) of the disk. By default, this column is hidden.
<b>Model</b>	Displays the model number of the disk. By default, this column is hidden.
<b>Vendor</b>	Displays the name of the disk vendor. By default, this column is hidden.
<b>Shelf ID</b>	Displays the ID of the shelf where the disk is located. By default, this column is hidden.
<b>Bay</b>	Displays the ID of the bay where the disk is located. By default, this column is hidden.

**Filters** Enables you to filter the disks based on the state of the disk. For example, you can use the filters to select spare and unassigned to view the list of all spare disks and unassigned disks in that cluster.

### Related Devices pane

The Related Devices pane enables you to view and navigate to the nodes, Vservers, 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 capacity of all the assigned disks minus the capacity of the broken disks.

**Vservers** Displays the number of the Vservers 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.

## 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 102
- [Capacity tab](#) on page 103
- [Disk Information tab](#) on page 104
- [Configuration tab](#) on page 105
- [History area](#) on page 106
- [Events list](#) on page 106
- [Related Devices pane](#) on page 106
- [Related Alerts pane](#) on page 106

### 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:

**Capacity** Displays the data capacity graph and the Snapshot copies graph, which display capacity details about the aggregate.

- **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.

**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, or RAID-DP).
- **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**

- **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, or RAID-DP).
- **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 and growth rates of the selected aggregate. You can select the graph you want to view from the drop-down list. You can also view details for a specified time period, such as one week, one month, or one year.

## 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** 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:

- Nodes** Displays the name, capacity, and the health status of the node which contains the aggregate. The capacity indicates the total capacity of all the assigned disks minus the capacity of the broken disks.
- 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

*[Performing suggested remedial actions for a full volume](#)* on page 56

## Job details page

The Job details page enables you to view status and other information about specific partner application 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.

<b>Duration</b>	Displays the elapsed time since the task began.
<b>Completed Time</b>	Displays the time the task completed. By default, this column is hidden.
<b>Task ID</b>	Displays the GUID that identifies an individual task for a job. The column can be sorted and filtered. By default, this column is hidden.
<b>Dependency order</b>	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.
<b>Task Details pane</b>	Displays additional information about each job task, including the task name, task description, and, if the task failed, a reason for the failure.
<b>Task Messages pane</b>	Displays messages specific to the selected task. Messages might include a reason for the error and suggestions for resolving it. Not all tasks display task messages.

## Definitions of user roles 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:

<b>Operator</b>	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.
<b>Storage Administrator</b>	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.
<b>OnCommand Administrator</b>	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.

## 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:

<b>Maintenance user</b>	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.
<b>Local user</b>	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.
<b>Remote group</b>	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.
<b>Remote user</b>	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.
<b>Database user</b>	Has read-only access to data in the Unified Manager database, has no access to the Unified Manager web interface or the maintenance console, and cannot execute API calls.

## Unified Manager 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	Operator	Storage Administrator	OnCommand Administrator
View storage system information	•	•	•
View other data like histories, capacity trends, etc.	•	•	•
View, assign, resolve events	•	•	•
View storage service objects, such as Vserver associations and resource pools	•	•	•
Manage storage service objects, such as Vserver associations and resource pools		•	•

Function	Operator	Storage Administrator	OnCommand Administrator
Define alerts		•	•
Manage storage management options		•	•
Manage storage management policies		•	•
Manage users			•
Manage administrative options			•
Manage database access			•

**Related references**

*Definitions of user types* on page 108

*Definitions of user roles in Unified Manager* on page 108

## Using the maintenance console

---

You can use the maintenance console to configure network settings, to configure and manage your virtual appliance, and to view server status to prevent and troubleshoot possible issues.

### Related concepts

*What the maintenance console does* on page 111

*Diagnostic user capabilities* on page 112

### Related tasks

*Sending a support bundle to technical support*

## What the maintenance console does

The maintenance console enables you to maintain the settings on your virtual appliance and to make any necessary changes to prevent issues from occurring.

You can use the maintenance console to perform the following actions:

- Troubleshoot any issues with your virtual appliance, especially if the Unified Manager web interface is not available
- Upgrade to newer versions of Unified Manager.
- Send AutoSupport messages to technical support
- Enable or disable remote access to the diagnostic shell
- Configure network settings
- Change the maintenance user password

### Related tasks

*Using the maintenance console* on page 111

## 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 OnCommand 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
- Enable or disable remote access to the diagnostic shell (only from VMware console)
- Generate support bundles to send to technical support

#### Related tasks

[Using the maintenance console](#) on page 111

## Diagnostic user capabilities

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

The diagnostic user can execute OS-level commands when directed by technical support, for troubleshooting purposes.

#### Related tasks

[Using the maintenance console](#) on page 111

## 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 have the maintenance user role.

#### 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.



### Related tasks

[Using the maintenance console](#) on page 111

## 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.

### Related tasks

[Using the maintenance console](#) on page 111

## Maintenance console menu

The maintenance console consists of different menus that enable you to maintain and manage your virtual appliance.

The maintenance console consists of the following menus:

- Upgrade OnCommand Unified Manager
- Network Configuration
- System Configuration
- Support/ Diagnostics

## Network Configuration menu

The Network Configuration menu enables you to manage the network settings. You should use this menu when the OnCommand Unified Manager user interface is not available.

The following menu choices are available.

<b>Display IP Address Settings</b>	Displays the current network settings for the virtual appliance, including the IP address, network, broadcast address, netmask, gateway, and DNS servers.								
<b>Change IP Address Settings</b>	Enables you to change any of the network settings for the virtual appliance, including the IP address, netmask, gateway, or DNS servers. If you switch your network settings from DHCP to static networking using the maintenance console, you cannot edit the host name. The host name provided by DHCP is used. Therefore, it is recommended to use the web UI. You must select <b>Commit Changes</b> for the changes to take place.								
<b>Display Domain Name Search Settings</b>	Displays the domain name search list used for resolving host names.								
<b>Change Domain Name Search Settings</b>	Enables you to change the domain names for which you want to search when resolving host names. You must select <b>Commit Changes</b> for the changes to take place.								
<b>Display Static Routes</b>	Displays the current static network routes.								
<b>Change Static Routes</b>	Enables you to add or delete static network routes. You must select <b>Commit Changes</b> for the changes to take place.								
	<table> <tr> <td><b>Add Route</b></td> <td>Enables you to add a static route.</td> </tr> <tr> <td><b>Delete Route</b></td> <td>Enables you to delete a static route.</td> </tr> <tr> <td><b>Back</b></td> <td>Takes you back to the <b>Main Menu</b>.</td> </tr> <tr> <td><b>Exit</b></td> <td>Exits the maintenance console.</td> </tr> </table>	<b>Add Route</b>	Enables you to add a static route.	<b>Delete Route</b>	Enables you to delete a static route.	<b>Back</b>	Takes you back to the <b>Main Menu</b> .	<b>Exit</b>	Exits the maintenance console.
<b>Add Route</b>	Enables you to add a static route.								
<b>Delete Route</b>	Enables you to delete a static route.								
<b>Back</b>	Takes you back to the <b>Main Menu</b> .								
<b>Exit</b>	Exits the maintenance console.								
<b>Disable Network Interface</b>	Disables any available network interfaces. If only one network interface is available, you cannot disable it. You must select <b>Commit Changes</b> for the changes to take place.								
<b>Enable Network Interface</b>	Enables available network interfaces. You must select <b>Commit Changes</b> for the changes to take place.								
<b>Commit Changes</b>	Applies any changes made to the network settings for the virtual appliance. You must select this option to enact any changes made, or the changes do not occur.								
<b>Ping a Host</b>	Pings a target host to confirm IP address changes or DNS configurations.								

<b>Restore to Default Settings</b>	Resets all settings to the factory default. You must select <b>Commit Changes</b> for the changes to take place.
<b>Back</b>	Takes you back to the <b>Main Menu</b> .
<b>Exit</b>	Exits the maintenance console.

## System Configuration menu

The System Configuration menu enables you to manage your virtual appliance, including viewing the server status, and rebooting and shutting down the virtual machine. You should use this menu when the OnCommand Unified Manager user interface is not available.

The following menu choices are available:

<b>Display Server Status</b>	Displays the current server status. Status options include Running and Not Running. If the server is not running, you might need to contact support.
<b>Reboot Virtual Machine</b>	Reboots the virtual machine, stopping all services. After rebooting, the virtual machine and services restart.
<b>Shut Down Virtual Machine</b>	Shuts down the virtual machine, stopping all services. The virtual machine does not restart. You can only select this option from the virtual machine console.
<b>Change &lt;logged in user&gt; User Password</b>	Enables you to change the password of the user currently logged in, which can only be the maintenance user.
<b>Increase Data Disk Size</b>	Enables you to increase the size of your data disks in the virtual machine.
<b>Increase Swap Disk Size</b>	Enables you to increase the size of your swap disks in the virtual machine.
<b>Change Time Zone</b>	Enables you to change the time zone to on your location.
<b>Change NTP Server</b>	Enables you to change the NTP Server settings such as IP address or FQDN.
<b>Back</b>	Returns you to the <b>Main Menu</b> .
<b>Exit</b>	Exits the maintenance console menu.

## Support and Diagnostics menu

The Support and Diagnostics menu enables you to manually send an AutoSupport message to technical support. You can also enable remote access through Secure Shell so that technical support personnel can assist you with troubleshooting issues.

The following menu choices are available:

<b>AutoSupport Submission</b>	Enables you to request that an AutoSupport message be generated and sent to technical support or other email recipients.
<b>Post to NetApp</b>	Posts the AutoSupport message to the NetApp AutoSupport web server.
<b>Send as email</b>	Sends the AutoSupport message via email to users who need a copy of the data, or to technical support.
<b>Both</b>	Sends the AutoSupport message to both the AutoSupport web server and via email to one or more recipients.
<b>Back</b>	Takes you back to the Main Menu.
<b>Exit</b>	Exits the maintenance console menu.
<b>Diagnostic Shell (Must be run on virtual machine console)</b>	Provides access to a diagnostic shell through the virtual machine console. This option should be used only when directed by technical support to diagnose issues, and not for normal maintenance activities. <b>Note:</b> After 15 minutes of inactivity, the console closes.
<b>Enable or Disable Remote Diagnostic Access (Must be run on virtual machine console)</b>	Enables or disables remote Secure Shell access to the diagnostic shell. You can only select this option from the virtual machine console. You must create a password to use when accessing the diagnostic shell through Secure Shell. This access expires automatically at midnight UTC time the following day and cannot be used again unless reactivated by the maintenance user. This option should be used only when directed by technical support to diagnose issues, and not for normal maintenance activities.
<b>Generate Support Bundle</b>	Enables you to create a 7-Zip file containing full diagnostic information in the diagnostic user's home directory. The file includes information generated by an AutoSupport message, the contents of the OnCommand Unified Manager database, detailed data about the Unified Manager server internals, and verbose-level logs not normally included in AutoSupport messages.

## Adding additional network interfaces

You can add new network interfaces if you need to separate network traffic.

### Before you begin

You must have added the network interface to the virtual appliance using vSphere.

The virtual appliance must be powered on.

### Steps

1. In the vSphere console **Main Menu**, select **System Configuration > Reboot Operating System**.

After rebooting, the maintenance console can detect the newly added network interface.

2. Access the maintenance console.
3. Select **Network Configuration > Enable Network Interface**.
4. Select the new network interface and press **Enter**.

### Example

Select **eth1** and press **Enter**.

5. Type **y** to enable the network interface.
6. Enter the network settings.

You are prompted to enter the network settings if using a static interface, or if DHCP is not detected.

After entering the network settings, you automatically return to the **Network Configuration** menu.

7. Select **Commit Changes**.

You must commit the changes to add the network interface.

### Related tasks

*[Accessing the maintenance console using the vSphere VM console](#) on page 113*

## Troubleshooting Unified Manager issues

---

If you encounter unexpected behavior during installation or when using Unified Manager, you can use specific troubleshooting procedures to identify and resolve the cause of such issues.

### VMware vSphere showing that VMware Tools are out-of-date

When you deploy the Unified Manager virtual appliance, the version of VMware Tools specific to your VMware environment is installed onto the virtual machine. If the virtual machine is booted on a newer version of VMware vSphere ESX, then VMware vSphere shows VMwareTools as out-of-date.

#### Workaround

Upgrade VMware Tools to the version specific to the new version of VMware vSphere ESX.

### Remote User option does not display in the Add User dialog box

**Issue** When a user opens the Add User dialog box, an alert displays indicating that remote authentication is not enabled and the Remote User and Remote Group options do not display in the Type drop-down list.

**Cause** Remote authentication has not been enabled in Unified Manager.

**Corrective Action** Enable remote authentication:

1. From the **Administration** menu, select **Setup Options**.
2. Open the **Management Server** list and select **Authentication**.
3. Select **Enable Remote Authentication**.
4. (Optional) You might also need to select an authentication service and add authentication servers, if those values are not already defined in the Authentication pane.

### Alerts are not received by designated recipients

**Issue** Alerts have been configured, but designated recipients are not receiving notifications.

**Cause** A possible cause is that alert settings are incorrectly set. For example, resources might be excluded that should not be, the wrong events might be selected, the wrong user might be selected, or the alert has not been enabled.

Another possible cause is that the notification options have not been correctly set.

**Corrective actions** Verify alert settings:

1. From the **Administration** menu, select **Manage Alerts**.
2. Select the problematic alert and click **Edit** to open the Edit Alert dialog box.
3. Click **Name** and verify that the Alert State option is Enabled.
4. Verify that the Resources, Events, and Recipients options are properly configured.

Verify notification settings:

1. From the **Administration** menu, select **Setup Options**.
2. Open the **General Settings** and select **Notification**.
3. Verify that a correct email address is entered in the From Address field.
4. If your environment requires SMTP for sending email, verify that the required information is entered.
5. If your network configuration requires SNMP, verify that the required information is entered.

## Glossary

---

### A

<b>Access Control List (ACL)</b>	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.
<b>admin Vserver</b>	In Clustered Data ONTAP, a Vserver that has overall administrative access to all objects in the cluster, including all objects owned by other Vservers, but does not provide data access to clients or hosts.
<b>aggregate</b>	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.
<b>aggregate committed capacity</b>	The data storage space in an aggregate that is committed to provide for its underlying volumes. Calculated by the total capacity provisioned for volumes.
<b>aggregate total capacity</b>	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.
<b>alert</b>	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. Alarms are used to monitor and manage datasets and resources. See also <i>event</i> and <i>severity type</i> .
<b>AutoSupport</b>	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.
<b>available capacity</b>	The amount of usable space available in a storage system. Calculated by the used capacity minus the unused reserve capacity.

### B

<b>backup relationship</b>	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 Clustered 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*.

**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, a group of connected nodes (storage systems) that share a global namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.

**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** Previous name for a *data Vserver*. See *data Vserver*.

**container object** An object such as an aggregate or a 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

<b>data capacity</b>	The storage space that is set aside by the container, such as aggregate or 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.
<b>Data ONTAP</b>	The operating system software running on NetApp storage devices.
<b>datastore</b>	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.
<b>data object</b>	A container of data such as a file, a directory, a volume, a LUN, that can be discovered, monitored, protected, created, or restored by the Unified Manager server.
<b>data Vserver</b>	In Clustered Data ONTAP, a virtual storage server that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data Vservers within a cluster.
<b>deduplication</b>	The consolidation of blocks of duplicate data into single blocks to store more information using less storage space.
<b>deduplication return</b>	The capacity savings resulting from deduplication. Calculated by the volume capacity before deduplication - the volume capacity after deduplication.
<b>destination</b>	The storage to which source data is backed up, mirrored, or migrated.
<b>destination data object</b>	A data object that contains the backed up or mirrored replicated data.
<b>dedupe</b>	See <i>deduplication</i> .
<b>DHCP</b>	See <i>Dynamic Host Configuration Protocol (DHCP)</i> .
<b>Dynamic Host Configuration Protocol (DHCP)</b>	The protocol for automating the assignment of network addresses.

## E

<b>ESX server</b>	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.

**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 Clustered Data ONTAP, 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.

**HA pair** In Clustered Data ONTAP, a pair of nodes (storage systems) configured to serve data for each other if one of the two nodes stops functioning.

**host** A computer system that accesses data on a storage system.

<b>host system</b>	A computer that accesses storage on a storage system.
<b>host bus adapter (HBA)</b>	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

<b>igroup</b>	initiator group. A collection of unique iSCSI node names of initiators (hosts) in an IP network that are given access to <i>front-end LUNs</i> when they are mapped to those LUNs. (Array LUNs on a storage array that provide storage for V-Series systems can be considered <i>back-end LUNs</i> .)
<b>incident</b>	Issues that have already impacted the availability or capacity of storage objects.
<b>incremental transfer</b>	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.
<b>initiator</b>	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.
<b>inode</b>	A data structure containing information about files on a storage system and in a UNIX file system.
<b>iSCSI</b>	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.
<b>iSCSI router</b>	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

<b>JBOD</b>	Just a Bunch Of Disks. An array of disks without any redundancy; that is, without RAID configuration.
<b>job</b>	Typically, a long-running operation. Some of the jobs include scheduled local backup of a dataset, a mirror transfer job, and password updates.

## K

**KB** Knowledge Base.

## 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. Unified Manager can be configured to point to an LDAP server for authentication of user requests; later versions of Data ONTAP can use Microsoft Active Directory, which uses LDAP.

**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 objects** Object types that represent storage containers such as volumes, qtrees, LUNs, and datasets.

**lower thresholds** The type of threshold that is set for generating an event when the counter value falls and remains below the lower threshold value for longer than the specified threshold interval.

## M

<b>maintenance user</b>	The user who has access rights to deploy and configure an OnCommand Unified Manager virtual appliance.
<b>Management Information Based (MIB)</b>	ASCII file that describes the information that the SNMP agent sends to network management stations.
<b>member</b>	Any data object that subscribes to or is created by a storage service.
<b>mirror (v)</b>	The process of creating an exact duplicate of all volume data from a source storage system to a destination storage system. In a NetApp configuration, mirror protection is based on SnapMirror technology.
<b>mirror copy (n)</b>	The exact duplicate of all volume data (both active and protected) from a NetApp source storage system to a destination storage system, created using NetApp Volume SnapMirror technology.
<b>move (v)</b>	To physically move data and any needed associated configuration of an object (for example, volume) from one aggregate to another within a cluster, including within a single node.

## N

<b>namespace</b>	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 <i>global namespace</i> .
<b>NDMP</b>	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.
<b>Network File System (NFS) export</b>	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.
<b>node</b>	In Clustered Data ONTAP, a storage system in a cluster. To distinguish between the two nodes in a high availability configuration, 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. The offline state also specifies the state of a storage object (for example, volume offline).

**online** A database state indicating that the database is available to users. The online state also specifies the state of a storage object (for example, volume online).

**OnCommand administrator** An RBAC role that allows 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 allows 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 in order 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.

**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** Issues that can impact the availability or capacity of storage objects.

**root member** A data object that subscribes to a storage service.



**S**

<b>SAN</b>	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.
<b>severity type</b>	Identifies the level of importance associated with an event. The severity type helps you determine priorities for taking corrective action.
<b>SFO</b>	See <i>storage failover (SFO)</i> .
<b>share</b>	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 <i>CIFS share</i> .
<b>Snapshot available capacity</b>	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.
<b>Snapshot copy</b>	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.
<b>Snapshot overflow</b>	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.
<b>Snapshot reserved capacity</b>	The storage space reserved for Snapshot copies.
<b>Snapshot return</b>	The storage space savings of Snapshot copies when compared with full volume copies. Calculated by subtracting the Snapshot capacity from the volume capacity.
<b>Snapshot reserve capacity</b>	The storage space that is set aside by the volume or the aggregate for its Snapshot copies. Data cannot be written to this space.
<b>Snapshot unused capacity</b>	The Snapshot reserve space remaining after the Snapshot copies are created.
<b>Snapshot used capacity</b>	The storage space used by the Snapshot copies in a volume or an aggregate.
<b>spare disk</b>	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.
<b>storable capacity</b>	The disk capacity that is available for use after right-sizing.

<b>standard HA configuration</b>	A configuration set up so that one node automatically takes over for its partner when the partner node becomes impaired.
<b>storage administrator</b>	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.
<b>storage controller</b>	The component of a storage system that runs the operating system and controls its disk subsystem. Storage controllers are also sometimes called <i>controllers</i> , <i>storage appliances</i> , <i>appliances</i> , <i>storage engines</i> , <i>heads</i> , <i>CPU modules</i> , or <i>controller modules</i> .
<b>storage efficiency</b>	The ratio of usable capacity to effective used capacity, accounting for efficiency returns. Calculated by the effective used capacity / the usable capacity.
<b>storage failover (SFO)</b>	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 <i>controller failover (CFO)</i> or <i>cluster failover (CFO)</i> .
<b>storage utilization</b>	The ratio of usable capacity to used capacity, without accounting for efficiency returns. Calculated by the used capacity / the usable capacity.
<b>system reserve capacity</b>	The capacity required for fixed system reserves, RAID parity, mirroring, and spare drives. Calculated by the fixed reserve + RAID reserve + spares.

## T

<b>takeover</b>	The emulation of the failed node identity by the takeover node in a high-availability configuration; the opposite of <i>giveback</i> .
<b>takeover node</b>	The functioning node (storage system) in an HA pair that manages access to the disk shelves and network connections of a failed partner node.
<b>thin provisioning</b>	A method of optimizing the efficiency with which the available space is used in storage, which is done by implementing on-demand allocation of disk storage space among multiple consumers, based on the minimum space required by each at any given time.
<b>trap</b>	An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the storage system.

## U

<b>unassigned disks</b>	Disks that are not assigned to any node and are not accounted as spare disks.
-------------------------	---

<b>unused capacity</b>	The usable storage space available for storing user data on a device. Also known as <i>available capacity</i> .
<b>unused reserve capacity</b>	The storage space allocated but unused by the aggregate Snapshot reserve, volume Snapshot reserve, volume fractional reserve, and Vol/LUN/File guaranteed space. These reserves can be adjusted by the user. Calculated by the aggregate Snapshot unused reserve + volume Snapshot unused reserve + volume fractional unused reserve + vol/LUN/file unused guaranteed space.
<b>usable capacity</b>	The storage space available to applications and users. Calculated by the raw capacity minus the system reserve capacity.
<b>used capacity</b>	The storage space used by applications or user data, including volume Snapshot copies and aggregate Snapshot copies. Calculated by the usable capacity minus the free capacity.

## V

<b>virtual appliance</b>	A prebuilt software solution containing virtual machines and software applications that are integrated, managed, and updated as a package. Also called <i>vApp</i> .
<b>vApp</b>	See <i>virtual appliance</i> .
<b>virtual machine (VM)</b>	A guest operating system and any application installed thereon, either running on a computing device on which OnCommand Unified Manager is installed or suspended to disk or any other storage media accessible by the computing device.
<b>virtual storage server</b>	See <i>Vserver</i> .
<b>VMware VirtualCenter (VC)</b>	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. A VC consists of agents, servers, and clients.
<b>volume</b>	A logical entity that contains user data that is accessible through one or more of the supported access protocols, including Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Fibre Channel (FC), and Internet SCSI (iSCSI).
<b>volume committed capacity</b>	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.

<b>volume total capacity</b>	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.
<b>Vserver</b>	In clustered Data ONTAP, a virtual storage server 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.
<b>Vserver guaranteed available capacity</b>	The data storage space that is guaranteed by the Vserver to its underlying volumes but is not used. Calculated as the sum of the available size of all the thick provisioned volumes.
<b>Vserver used capacity</b>	The data storage space that is guaranteed by the Vserver to its underlying volumes. Calculated as the sum of the used data capacity of all the volumes associated with the Vserver.
<b>Vserver unguaranteed capacity</b>	The data storage space that is not guaranteed by the Vserver to its underlying volumes. Calculated as the sum of the available sizes of all the thin provisioned volumes.
<b>Vserver total capacity</b>	The sum of the total data storage space of all the volumes in the Vserver.
<b>Vserver unguaranteed capacity</b>	The data storage space that is not guaranteed by the Vserver to its underlying volumes. Calculated as the Vserver total capacity minus the sum of the committed capacity of aggregates that are associated with the Vserver.

## Copyright information

---

Copyright © 1994–2013 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, Mars, Manage ONTAP, 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](http://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](http://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](mailto: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

6.0 release of Unified Manager  
related concepts [7](#)

## A

acknowledging  
events [68](#)

adding  
alerts [31, 50](#)  
authentication servers [27](#)  
clusters [33, 46](#)  
new rules [52](#)  
remote groups [30](#)  
remote users [30](#)  
rules [52](#)

administrative tasks  
common workflows for performing [19](#)

administrators  
OnCommand [108](#)  
storage [108](#)

Aggregate details page [102](#)

aggregates  
configuring global threshold values for [29](#)  
details about [102](#)

alerts  
adding [31, 50](#)  
configuring your environment for [25](#)  
creating [31, 50](#)

alerts not received  
troubleshooting [118](#)

assigning  
events [68](#)  
user roles [30](#)

authentication  
adding servers [27](#)  
enabling remote [26](#)

AutoSupport  
using the maintenance console [115](#)

availability events  
general description [13](#)

availability health  
defined [13](#)

availability issues  
correcting a flash card offline condition [35](#)

correcting a storage failover interconnect link down  
condition [37, 38](#)  
resolving volume offline conditions [40, 41](#)

availability workflows  
introduction [34](#)

## B

backup process  
overview [16](#)

backup vault protection  
configuration [15](#)

backup vault protection relationships  
defined [14](#)

## C

capabilities  
database user [13](#)  
FlexVol volume [10](#)  
table of roles associated with [109](#)

capacity  
information for volumes [73](#)

capacity events  
full volume [56](#)  
resolving [55](#)  
suggested remedial actions for a full volume  
list of [56](#)  
performing [56](#)

certificates  
generating HTTPS security certificates [23](#)  
viewing HTTPS security certificates [23](#)

cluster  
description of [7](#)  
epsilon [7](#)  
quorum [7](#)

Cluster details page [94](#)

clustered Data ONTAP systems  
adding [33, 46](#)

clusters  
adding [33, 46](#)  
details about [94](#)

configuring  
aggregate global threshold values [29](#)  
DNS [22](#)  
network settings [22](#)



- notification settings [26](#)
- thresholds [28](#)
- volume global threshold values [29](#)
- your environment [20](#)
- creating
  - alerts [31](#), [50](#)
  - custom rules [53](#)
  - database users [30](#)
  - local users [30](#)
  - rules [52](#)
- custom rules
  - creating [53](#)

## D

- data disk size
  - increasing [115](#)
- data policies
  - configuration [83](#)
  - defined [12](#)
  - exporting [54](#)
- Data Policy tab [83](#)
- database users
  - capabilities [13](#)
  - creating [30](#)
  - defined [108](#)
- DHCP
  - enabling [22](#)
- diag users [112](#)
- DNS
  - configuring [22](#)

## E

- editing
  - Infinite Volume threshold settings [47](#)
  - network settings [22](#)
  - storage class threshold settings [49](#)
  - unmanaged relationship lag threshold settings [30](#)
- efficiency
  - information for volumes [73](#)
- enabling
  - DHCP [22](#)
- environment
  - setup [20](#)
- Event details page [69](#)
- event impact areas
  - availability [72](#)
  - capacity [72](#)
  - configuration [72](#)

- description [72](#)
- protection [72](#)
- event impact levels
  - description [72](#)
  - event [72](#)
  - incident [72](#)
  - risk [72](#)
- event severity types
  - critical [71](#)
  - description [71](#)
  - error [71](#)
  - information [71](#)
  - warning [71](#)
- event states
  - definitions [17](#)
- events
  - acknowledged [17](#)
  - acknowledging [68](#)
  - adding notes about [67](#)
  - assigning to users [68](#)
  - definition of [17](#)
  - details [69](#)
  - impact areas [72](#)
  - impact levels [72](#)
  - new [17](#)
  - obsolete [17](#)
  - resolved [17](#)
  - resolving [68](#)
  - reviewing notes about [67](#)
  - severity types [71](#)
  - states [17](#)
  - viewing notes about [67](#)

## F

- failed protection jobs
  - identifying [57](#)
  - identifying the cause [58](#)
  - performing corrective actions [58](#)
  - resolving [57](#)
- failure
  - protection job [57](#)
- flash card offline condition
  - troubleshooting [35](#)
- FlexVol volumes
  - capabilities of [10](#)

## G

- generation of a support bundle

purpose [64](#)

## H

host name

changing [21](#)

HTTPS

generating new security certificates [23](#)

viewing the security certificate [23](#)

## I

impact areas

availability [72](#)

capacity [72](#)

configuration [72](#)

description [72](#)

protection [72](#)

impact levels

description [72](#)

event [72](#)

incident [72](#)

risk [72](#)

Infinite Volumes

definition of [10](#)

editing threshold settings [47](#)

storage classes, definition of [11](#)

workflow for managing [47](#)

workflow for monitoring [45](#)

workflow for setting up [45](#)

issues

workflows for troubleshooting common [19](#)

## J

Job Details page [107](#)

jobs

defined [11](#)

identifying cause of failure [57](#)

list of those you can monitor [11](#)

progress monitoring [107](#)

resolving terminated [57](#)

status [107](#)

troubleshooting failures [107](#)

## L

lag issues

resolving [61](#)

local users

creating [30](#)

defined [108](#)

## M

maintenance console

accessing remote diagnostic shell [115](#)

accessing using Secure Shell [64](#), [112](#)

accessing using VM console [113](#)

AutoSupport [115](#)

diag users

capabilities [112](#)

generating a support bundle [64](#)

overview [6](#)

purpose [111](#)

restarting the virtual machine [24](#)

restarting Unified Manager [24](#)

role of maintenance user [111](#)

support and diagnostics [115](#)

system configuration [115](#)

what it does [111](#)

maintenance user

defined [108](#)

what it does [111](#)

managing

Infinite Volumes, workflow for [47](#)

mirror protection

configuration [15](#)

mirror protection relationships

defined [14](#)

modifying

Infinite Volume threshold settings [47](#)

storage class threshold settings [49](#)

unmanaged relationship lag threshold settings [30](#)

monitoring

Infinite Volumes, workflow for [45](#)

protection relationships [57](#)

## N

namespaces

in Vservers [8](#)

Network Configuration

using the maintenance console [114](#)

network interfaces

adding new [116](#)

network settings

configuring [22](#)

customizing the host name [21](#)

- editing [22](#)
- notification
  - adding alerts [31](#), [50](#)
  - configuring settings [26](#)

## O

- offline flash card condition
  - troubleshooting [35](#)
- offline volume
  - determining cause [42](#), [44](#)
- offline volume condition
  - troubleshooting [40](#)
- offline volumes
  - determining cause [43](#)
- OnCommand administrators
  - defined [108](#)
- operators
  - defined [108](#)

## P

- passwords
  - changing [34](#)
- physical storage
  - adding clusters [33](#), [46](#)
- policies
  - exporting data policies [54](#)
- protection job failure
  - identifying [57](#)
  - resolving [57](#)
- protection job failures
  - identifying the cause [58](#)
  - performing corrective actions [58](#)
- protection jobs
  - correcting failed [58](#)
  - execution [15](#)
  - identifying failed [57](#)
  - resolving failed [57](#)
- protection relationships
  - creation [15](#)
  - lag issue resolution workflow [61](#)
  - monitoring [57](#)
  - troubleshooting [57](#)
- purpose of maintenance console
  - list of actions performed using [111](#)

## Q

- Qtrees tab [83](#)

## R

- reassigning
  - events [68](#)
- reference information
  - common to workflows [67](#)
- relationships
  - unmanaged, editing lag thresholds settings for [30](#)
- releases of Unified Manager
  - concepts to understand for working with 6.0 [7](#)
- remote authentication
  - enabling [26](#)
- remote diagnostic shell
  - accessing through Secure Shell [115](#)
- remote groups
  - adding [30](#)
  - defined [108](#)
- Remote User UI option
  - troubleshooting lack of display [118](#)
- remote users
  - adding [30](#)
  - defined [108](#)
- resolving
  - events [68](#)
- resource pools
  - about [12](#)
- resources
  - selecting using Vserver associations [13](#)
- restore and recovery operations
  - tools for performing [16](#)
- restore process
  - overview [16](#)
- reviewing
  - notes about events [67](#)
- role-based access control
  - See* RBAC
- roles
  - assigning to users [30](#)
  - defined [108](#)
  - table of capabilities associated with [109](#)
- rules
  - adding [52](#)
  - creating [52](#)
  - creating, custom [52](#), [53](#)
  - creating, using templates [52](#)
  - defined [12](#)
  - exporting [54](#)

**S**

- Secure Shell
  - using to access the maintenance console [64](#), [112](#)
- security certificates
  - generating, HTTPS [23](#)
  - viewing, HTTPS [23](#)
- setting up
  - aggregate global threshold values [29](#)
  - notification settings [26](#)
  - SMTP server [26](#)
  - SNMP [26](#)
  - thresholds [28](#)
  - volume global threshold values [29](#)
- setup
  - post-deployment [20](#)
- severity types
  - critical [71](#)
  - description [71](#)
  - error [71](#)
  - information [71](#)
  - warning [71](#)
- SnapMirror license
  - enabling mirror protection [14](#)
- SnapVault license
  - enabling backup vault protection [14](#)
- storage administrators
  - defined [108](#)
- storage classes
  - capacity details of [83](#)
  - definition of [11](#)
  - editing threshold settings [49](#)
- storage failover link down condition
  - troubleshooting [37](#), [38](#)
- support bundles
  - generating [64](#)
  - retrieving using a Windows client [65](#)
  - retrieving using the CLI [65](#)
  - sending to technical support [63](#)
  - sending to technical support for diagnosis [65](#)
  - uploading to technical support [67](#)
- swap disk size
  - increasing [115](#)
- system configuration menu
  - changing user password [115](#)
  - displaying server status [115](#)
  - rebooting virtual machine [115](#)
  - shutting down virtual machine [115](#)
- System configuration menu
  - using the maintenance console [114](#)

**T**

- tasks
  - common to workflows [67](#)
  - viewing information about [107](#)
- technical support
  - sending a support bundle to [63](#)
- thresholds
  - configuring [28](#)
  - editing settings for Infinite Volumes [47](#)
  - editing settings for storage classes [49](#)
  - editing settings for unmanaged relationships [30](#)
  - global values for aggregates [29](#)
  - global values for volumes [29](#)
- time zone
  - changing [115](#)
- troubleshooting
  - alerts not received by designated recipients [118](#)
  - flash card offline condition [35](#)
  - job failures [107](#)
  - procedures for [118](#)
  - protection relationships [57](#)
  - Remote User option does not display [118](#)
  - sending a support bundle [63](#)
  - storage failover link down condition [37](#), [38](#)
  - VMware Tools [118](#)
  - volume offline condition [40](#)
- types
  - of users [108](#)
- types of users
  - database users [13](#)
  - maintenance user [111](#)

**U**

- Unified Manager web UI
  - overview [6](#)
- unmanaged relationships
  - editing lag thresholds settings for [30](#)
- user roles
  - assigning [30](#)
- users
  - adding [30](#)
  - capabilities associated with [109](#)
  - changing passwords [34](#)
  - creating [30](#)
  - database [13](#)
  - diagnostic user [112](#)
  - maintenance user [111](#)
  - roles [108](#)

types [108](#)

## V

viewing

notes about events [67](#)

virtual machine

restarting [24](#)

virtual machine console

accessing the maintenance console [113](#)

virtual machines

changing maintenance user password [115](#)

rebooting [115](#)

shutting down [115](#)

virtual servers

*See* Vservers

VM console

accessing the maintenance console [113](#)

VMware Tools

troubleshooting [118](#)

Volume details page [73](#)

volume offline

determining if caused by a down host cluster node  
[42](#)

determining if caused by a stopped Vserver resulting  
from a down cluster node [43](#)

determining if caused by broken RAID disks [44](#)

volume offline condition

troubleshooting [40](#)

volumes

capacity information [73](#)

configuring global threshold values for [29](#)

details about [73](#)

efficiency information [73](#)

how they work [9](#)

Infinite Volume defined [10](#)

provisioning using Vserver associations [13](#)

Vserver associations

about [13](#)

Vserver details page [83](#)

Vservers

defined [8](#)

details about [83](#)

namespaces [8](#)

## W

what the diagnostic user does

defined [112](#)

Windows client

retrieving the support bundle [65](#)

workflows

list of common [19](#)

reference information common to [67](#)

resolving lag issues [61](#)

summarized [19](#)

tasks common to [67](#)