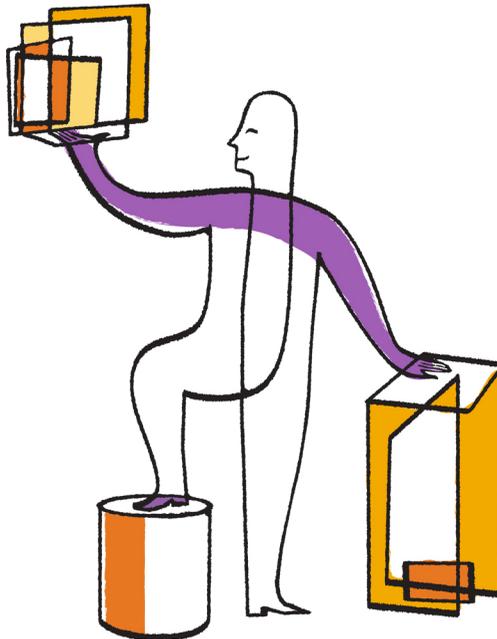




NetApp®

OnCommand® Unified Manager 6.1

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
Feedback: doccomments@netapp.com

Part number: 215-08262_A0
February 2014

Contents

Unified Manager product and Administration Guide overview	7
Concepts related to working with Unified Manager	8
What a cluster is	9
What SVMs are	9
How volumes work	11
What a FlexVol volume is	12
Capabilities that FlexVol volumes provide	12
What an Infinite Volume is	13
What a storage class is	13
What jobs are	13
What resource pools are	14
What rules and data policies are	14
Understanding SVM associations	15
Database user capabilities	15
What availability health is	15
What mirror and backup vault protection relationships are	16
How protection relationships are created and protection jobs run	17
Virtual appliance backup and restore process overview	18
What events are	18
Event state definitions	19
What annotations are	19
What performance incidents are	20
Display of performance incidents in Unified Manager	20
Workloads, cluster components, and performance	21
Purpose of a connection between Performance Manager and Unified Manager	23
Connections between multiple Performance Manager servers and Unified Manager	24
Common Unified Manager administrative workflows and tasks	25
Configuring your environment after deployment	26
Changing the Unified Manager host name	27
Configuring Unified Manager to send alert notifications	31
Adding clusters	40

Changing the local user password	41
Monitoring and troubleshooting data availability	42
Resolving a flash card offline condition	42
Scanning for and resolving storage failover interconnect link down conditions	44
Resolving volume offline issues	47
Diagnosing performance issues using Unified Manager and Performance Manager	53
Setup tasks for a connection between Performance Manager and Unified Manager	53
Creating a user with Event Publisher role privileges	53
Configuring a connection between a Performance Manager server and the Unified Manager server	54
Analyzing a performance incident	55
Setting up and monitoring an SVM with Infinite Volume without storage classes	58
Adding clusters	59
Editing the Infinite Volume threshold settings	60
Managing your Infinite Volume with storage classes and data policies	61
Editing the threshold settings of storage classes	63
Adding an alert	63
Creating rules	66
Exporting a data policy configuration	68
Resolving capacity issues	68
Performing suggested remedial actions for a full volume	69
Creating, monitoring, and troubleshooting protection relationships	70
Setting up protection relationships in Unified Manager	71
Performing a protection relationship failover and failback	76
Resolving a protection job failure	81
Resolving lag issues	85
Restoring data from Snapshot copies	87
Restoring data using the Volume details page	87
Restoring data using the Volumes page	88
Prioritizing storage object events using annotations	89
Creating rules to annotate storage objects	89
Viewing annotations	90

Removing annotations for storage objects	90
Understanding more about annotations	91
Sending a support bundle to technical support	91
Accessing the maintenance console using Secure Shell	92
Generating a support bundle	93
Retrieving the support bundle using a Windows client	93
Retrieving the support bundle using a UNIX or Linux client	94
Sending a support bundle to technical support	95
Related tasks and reference information	96
Adding and reviewing notes about an event	96
Assigning events	96
Acknowledging and resolving events	97
Event details page	98
Description of event severity types	101
Description of event impact levels	101
Description of event impact areas	101
Volume details page	102
Storage Virtual Machine details page	116
Cluster details page	131
Aggregate details page	140
Job details page	146
Definitions of user roles in Unified Manager	147
Definitions of user types	148
Unified Manager roles and capabilities	148
Using the maintenance console	150
What the maintenance console does	150
What the maintenance user does	150
Diagnostic user capabilities	151
Accessing the maintenance console using Secure Shell	151
Accessing the maintenance console using the vSphere VM console	152
Maintenance console menu	152
Network Configuration menu	152
System Configuration menu	154
Support and Diagnostics menu	154
Adding additional network interfaces	155
Troubleshooting Unified Manager issues	157

Incorrect trigger condition displayed for Aggregate Snapshot Reserve Full event	157
VMware vSphere showing that VMware Tools are out-of-date	157
Remote User option does not display in the Add User dialog box	157
Alerts are not received by designated recipients	158
Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring	159
Glossary	161
Copyright information	175
Trademark information	176
How to send your comments	177
Index	178

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.

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

Unified Manager web UI

The Unified Manager web UI enables a storage administrator, cluster administrator, or SVM administrator to monitor and troubleshoot cluster or Storage Virtual Machine (SVM, formerly known as Vserver) issues relating to data storage capacity, availability, performance, and protection.

This guide describes some common workflows that an administrator can follow to troubleshoot storage capacity, data availability, performance, 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 6.1 requires you to be familiar with several concepts, some of which are new in this release.

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

Related concepts

What a cluster is on page 9

What SVMs are on page 9

How volumes work on page 11

What a FlexVol volume is on page 12

Capabilities that FlexVol volumes provide on page 12

What an Infinite Volume is on page 13

What a storage class is on page 13

What jobs are on page 13

What resource pools are on page 14

What rules and data policies are on page 14

Understanding SVM associations on page 15

Database user capabilities on page 15

What availability health is on page 15

What mirror and backup vault protection relationships are on page 16

How protection relationships are created and protection jobs run on page 17

Virtual appliance backup and restore process overview on page 18

What events are on page 18

Event state definitions on page 19

What annotations are on page 19

What performance incidents are on page 20

Display of performance incidents in Unified Manager on page 20

Workloads, cluster components, and performance on page 21

Purpose of a connection between Performance Manager and Unified Manager on page 23

Connections between multiple Performance Manager servers and Unified Manager on page 24

Related references

Glossary on page 161

What a cluster is

You can group pairs of nodes together to form a scalable cluster. Creating a cluster enables the nodes to pool their resources and distribute work across the cluster, while presenting administrators with a single entity to manage. Clustering also enables continuous service to end users if individual nodes go offline.

The maximum number of nodes within a cluster depends on the platform model and licensed protocols. For details about cluster size limits, see the *Hardware Universe* at hwu.netapp.com.

Each node in the cluster can view and manage the same volumes as any other node in the cluster. The total file-system namespace, which comprises all of the volumes and their resultant paths, spans the cluster.

When new nodes are added to a cluster, there is no need to update clients to point to the new nodes. The existence of the new nodes is transparent to the clients.

If you have a two-node cluster, you must configure cluster (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 8

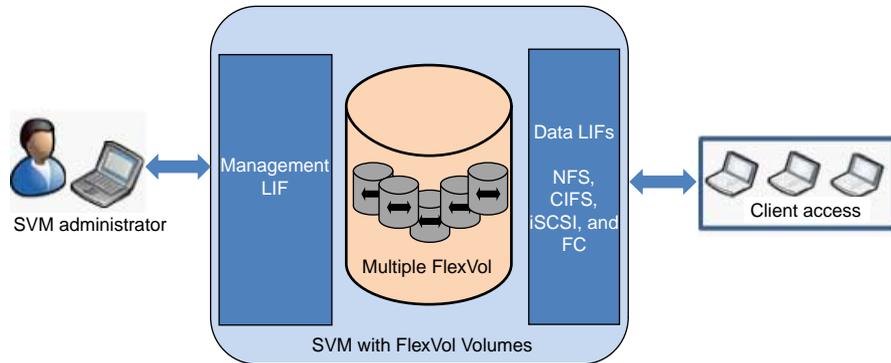
What SVMs are

Storage Virtual Machines (SVMs, formerly known as Vservers) contain data volumes and one or more LIFs through which they serve data to the clients. Starting with clustered Data ONTAP 8.1.1, SVMs can either contain one or more FlexVol volumes, or a single Infinite Volume.

SVMs securely isolate the shared virtualized data storage and network, and each SVM appears as a single dedicated server to the clients. Each SVM has a separate administrator authentication domain and can be managed independently by its SVM administrator.

A cluster can have one or more SVMs with FlexVol volumes and SVMs with Infinite Volume.

SVM with FlexVol volumes

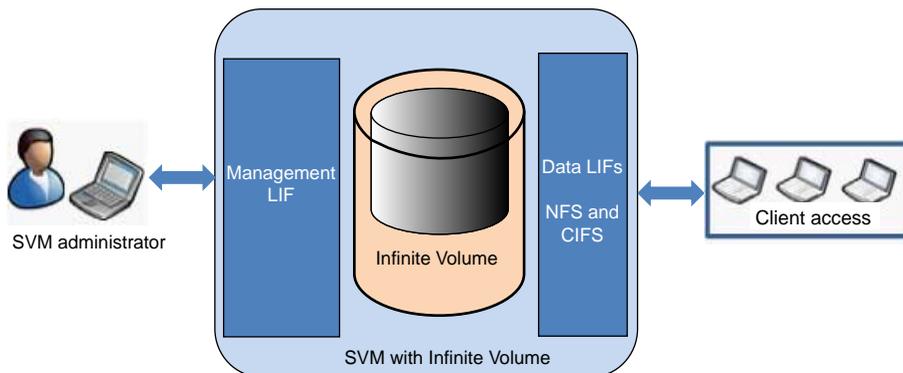


Each SVM with FlexVol volumes in a NAS environment presents a single directory hierarchical view and has a unique namespace. The namespace enables NAS clients to access data without specifying the physical location of the data. The namespace also enables the cluster and SVM administrators to manage distributed data storage as a single directory with multiple levels of hierarchy.

The volumes within each NAS SVM are related to each other through junctions and are mounted on junction paths. These junctions present the file system in each volume. The root volume of the SVM is a FlexVol volume that resides at the top level of the namespace hierarchy; additional volumes are mounted to the SVM root volume to extend the namespace. As volumes are created for the SVM, the root volume of the SVM contains junction paths.

SVMs with FlexVol volumes can contain files and LUNs. They provide file-level data access by using NFS and CIFS protocols for the NAS clients, and block-level data access by using iSCSI and Fibre Channel (FC) (FCoE included) for SAN hosts.

SVM with Infinite Volume



SVMs with Infinite Volume can contain only one Infinite Volume to serve data. Each SVM with Infinite Volume includes only one junction path, which has a default value of `/NS`. The junction provides a single mount point for the large namespace provided by the SVM with Infinite Volume. You cannot add more junctions to an SVM with Infinite Volume. However, you can increase the size of the Infinite Volume.

SVMs with Infinite Volume can contain only files. They provide file-level data access by using NFS and CIFS (SMB 1.0) protocols. SVMs with Infinite Volume cannot contain LUNs and do not provide block-level data access.

Note: The Data ONTAP command-line interface (CLI) continues to use the term `Vserver` in the output, and `vserver` as a command or parameter name has not changed.

Related concepts

[Concepts related to working with Unified Manager](#) on page 8

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 Storage Virtual Machine (SVM). The SVM is a virtual management entity, or server, that consolidates various cluster resources into a single manageable unit. When you create a volume, you specify the SVM it is associated with. The type of the volume (FlexVol volume or Infinite Volume) is determined by an immutable SVM attribute.

Volumes have a language. The language of the volume determines the character set Data ONTAP uses to display file names and data for that volume. The default value for the language of the volume is the language of the SVM.

Volumes depend on their associated aggregates for their physical storage; they are not directly associated with any concrete storage objects, such as disks or RAID groups. If the cluster administrator has assigned specific aggregates to an SVM, then only those aggregates can be used to provide storage to the volumes associated with that SVM. This impacts volume creation, and also copying and moving FlexVol volumes between aggregates.

For more information about Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

For more information about SVMs, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

For more information about data protection mirrors, see the *Clustered Data ONTAP Data Protection Guide*.

For more information about physical storage resources such as aggregates, disks, and RAID groups, see the *Clustered Data ONTAP Physical Storage Management Guide*.

Related concepts

[Concepts related to working with Unified Manager](#) on page 8

What a FlexVol volume is

A FlexVol volume is a data container associated with a Storage Virtual Machine (SVM) 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 8

What a storage class is

A storage class is a definition of aggregate characteristics and volume settings. You can define different storage classes and associate one or more storage classes with an Infinite Volume. You must use OnCommand Workflow Automation to define workflows for your storage class requirements and to assign storage classes to Infinite Volumes.

You can define the following characteristics for a storage class:

- Aggregate characteristics, such as the type of disks to use
- Volume settings, such as compression, deduplication, and volume guarantee

For example, you can define a storage class that uses only aggregates with SAS disks and the following volume settings: thin provisioning with compression and deduplication enabled.

Related concepts

Concepts related to working with Unified Manager on page 8

What jobs are

A job is a series of tasks that you can monitor using Unified Manager. Viewing jobs and their associated tasks enables you to determine if they have completed successfully.

Jobs are initiated when you create SnapMirror and SnapVault relationships, when you perform any relationship operation (break, edit, quiesce, remove, resume, resynchronize, and reverse resync), when you perform data restoration tasks, when you log in to a cluster, and so on.

When you initiate a job, you can use the Jobs page and the Job details page to monitor the job and the progress of the associated job tasks.

Related concepts

Concepts related to working with Unified Manager on page 8

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 8

What rules and data policies are

A *rule* determines the placement of files (data) in a Storage Virtual Machine (SVM) with Infinite Volume. A collection of such rules is known as a *data policy*.

Rule Rules mainly consist of a set of predefined conditions and information that determine where to place files in the Infinite Volume. When a file is placed in the Infinite Volume, the attributes of that file are matched with the list of rules. If attributes match the rules, then that rule's placement information determines the storage class where the file is placed. A default rule in the data policy is used to determine the placement of files if the attributes do not match any of the rules in the rule list.

For example, if you have a rule, "Place all files of type .mp3 in the bronze storage class.," all .mp3 files that are written to the Infinite Volume would be placed in the bronze storage class.

Data policy A data policy is a list of rules. Each SVM with Infinite Volume has its own data policy. Each file that is added to the Infinite Volume is compared to its data policy's rules to determine where to place that file. The data policy enables you to filter incoming files based on the file attributes and place these files in the appropriate storage classes.

Related concepts

Concepts related to working with Unified Manager on page 8

Understanding SVM associations

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

Associations are always created between the primary SVM and the destination regardless of whether the destination is a secondary or a tertiary destination. You cannot associate a secondary destination SVM with a tertiary destination SVM.

You can associate SVMs in two ways:

- **Associate any SVM**
You can create an association between any primary SVM source to one or more destination SVMs. This means that all existing SVM that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVMs. For example, you might want applications from several different sources at different locations backed up to one or more destination SVMs in one location.
- **Associate a particular SVM**
You can create an association of a specific source SVM with one or more destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, you can choose this option to associate a specific SVM source to a specific SVM destination that is assigned to only that client.

Related concepts

[Concepts related to working with Unified Manager](#) on page 8

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 8

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 8

Related tasks

Resolving a flash card offline condition on page 42

Scanning for and resolving storage failover interconnect link down conditions on page 44

Resolving volume offline issues on page 47

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 8

How protection relationships are created and protection jobs run

Unified Manager enables you to discover, monitor, troubleshoot, and manage event resolution for existing SnapMirror and SnapVault protection relationships, and it enables you to configure new protection relationships.

You can monitor and troubleshoot protection relationship issues from the Dashboard page.

You can create protection relationships with Unified Manager in the Protection view of the Volumes page, or from the Volume details page, but you first must install and configure OnCommand Workflow Automation and then pair the two applications so they work together.

OnCommand Workflow Automation provides a user interface with pre-configured workflows, which can be executed to configure SnapMirror or SnapVault protection. You can also use some third-party applications to configure protection relationships.

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>.

The following management tools can also enable cluster or Storage Virtual Machine (SVM) administrators to manage SnapMirror or SnapVault protection:

- OnCommand System Manager
OnCommand System Manager or later provides a web interface with guided prompts to manage SnapMirror and SnapVault protection.
For more information, see the *OnCommand System Manager Help*, which you can access from within OnCommand System Manager.
For information about the version of OnCommand System Manager you should use with Unified Manager, see Interoperability Matrix at support.netapp.com/matrix.
- Data ONTAP CLI commands
The Data ONTAP CLI provides commands to manage SnapMirror and SnapVault protection.
For more information about managing protection relationships through the Data ONTAP CLI, see the documentation about mirror and SnapVault backup 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 8

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

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 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.

Related concepts

[Concepts related to working with Unified Manager](#) on page 8

What events are

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

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

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

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

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

Related concepts

Concepts related to working with Unified Manager on page 8

Event state definitions

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

The different event states are as follows:

New The state of a new event.

Acknowledged The state of an event when you have acknowledged it.

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

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

Note: You cannot acknowledge or resolve an obsolete event.

Example of different states of an event

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

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

Related concepts

Concepts related to working with Unified Manager on page 8

What annotations are

Annotations enable you to dynamically tag storage objects using user-defined rules. When you tag storage objects using annotations, you can filter and view the events that are related to them. You can apply annotations to clusters, volumes, and Storage Virtual Machines (SVMs).

OnCommand Unified Manager generates several events based on the availability, capacity, protection, and performance of storage objects in your data center. Annotations help you to filter and identify only those events that correspond to the storage objects that are tagged.

Related concepts

[Concepts related to working with Unified Manager](#) on page 8

[Prioritizing storage object events using annotations](#) on page 89

What performance incidents are

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

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

Related concepts

[Concepts related to working with Unified Manager](#) on page 8

Related tasks

[Configuring a connection between a Performance Manager server and the Unified Manager server](#) on page 54

Display of performance incidents in Unified Manager

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

Where performance incidents are displayed

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

What information the performance graph displays

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

What information the performance incident listings display

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

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

What states of performance incidents are displayed

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

Related concepts

[*Concepts related to working with Unified Manager*](#) on page 8

Related tasks

[*Configuring a connection between a Performance Manager server and the Unified Manager server*](#) on page 54

Workloads, cluster components, and performance

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

What a workload is

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

Workloads that are monitored by Performance Manager include the following:

FlexVol volume workload	<p>The storage activity associated with maintaining a volume and its read/write activity.</p> <p>Multiple FlexVol volumes constitute multiple workloads that must be supported by the shared cluster components on which the volumes reside or through which volume data is throughput.</p> <p>Because FlexVol volumes are created and configured by storage administrators, volume workloads are also referred to as <i>user-defined workloads</i>.</p>
Deduplication workload	<p>The storage activity associated with data deduplication activity on a storage cluster. Because deduplication is a system-defined feature, deduplication workloads are referred to as <i>system workloads</i>.</p>
RAID reconstruction workload	<p>The storage activity associated with RAID reconstruction of stored data that might be temporarily lost due to a single or double disk failure in a storage array. RAID reconstruction workloads are system workloads.</p>
Snapshot copy workload	<p>The storage activity associated with the automatic backup imaging of data in Snapshot copies. Snapshot copy workloads are system workloads.</p>
WAFL consistency point (CP) workload	<p>The storage activity associated with maintaining the WAFL consistency point feature, which enables fast restart and recovery of storage cluster nodes in the event of improper shutdown. WAFL consistency point workloads are system workloads.</p>
WAFL scan	<p>The storage activity associated with execution of the WAFL scan operation. WAFL scan operations are system workloads.</p>
Volume moves workload	<p>The storage activity associated with flexible volume migration. Volume move workloads are system workloads.</p>
Other internal activity	<p>The storage activity associated with other system-related operations.</p>

What a cluster component is

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

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

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

Network	The I/O processing from clients connected to the cluster. If network issues occur outside of the cluster (for example, high network traffic or over-utilized clients), thus causing slow workload I/O response times, Performance Manager generates a performance incident related to this.
Policy Group limit	The Storage Quality of Service (QoS) policy group of which the workload is a member. If the processing demand of the workloads in the policy group exceed the policy limit for I/O operations, the workloads are throttled, and Performance Manager generates a performance incident.
Network Processing	The software component in the cluster involved with I/O processing between the network protocols and the cluster. If high response time is detected during network processing, Performance Manager generates a performance incident with the name of the cluster node that is handling that network processing.
Cluster Interconnect	The cables and adapters with which clustered nodes are physically connected. Performance Manager generates a performance incident if it detects contention in the interconnect.
Data Processing	The software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains a volume workload. If Performance Manager detects a high response time between the cluster and the storage aggregate, it generates a performance incident with the name of the cluster node that is handling the data processing. The I/O requests might or might not be from other nodes in the cluster.
Aggregate	The storage aggregate that contains the volume workload. If the processing demands of one workload increase the response time of other workloads on the aggregate, Performance Manager generates a performance incident with the name of the aggregate in contention.

Related concepts

Concepts related to working with Unified Manager on page 8

Purpose of a connection between Performance Manager and Unified Manager

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

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

Related concepts

Concepts related to working with Unified Manager on page 8

Connections between multiple Performance Manager servers and Unified Manager

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

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

Related concepts

Concepts related to working with Unified Manager on page 8

Common Unified Manager administrative workflows and tasks

Some common administrative workflows and tasks associated with Unified Manager include selecting the storage clusters to monitor; diagnosing conditions that adversely affect data availability, performance, capacity, and protection; creation of protection relationships; restoring lost data; configuration and management of Infinite Volumes; and, when necessary, bundling and sending diagnostic data to technical support.

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

The most important cluster, Storage Virtual Machine (SVM), 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 administrative tasks are described in this document:

- **Setting up the management environment after deployment**
After storage clusters and their storage resources have been configured using the Data ONTAP CLI or System Manager, storage administrators 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, storage administrators can follow embedded links to display connectivity information about the affected storage resource, display troubleshooting advice, and assign issue resolution to other administrators.
- **Configuring and monitoring for performance incidents**
After setting up a connection between Performance Manager and Unified Manager, the OnCommand Administrator can monitor the performance of resources monitored by the two applications.
- **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, storage administrators can use Unified Manager to monitor, set notification thresholds, and define data policy for that volume and its storage classes. Optionally, storage administrators 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, storage administrators can follow embedded links to display storage capacity current and historical trending information about the affected volume, display troubleshooting advice, and assign issue resolution to other administrators.

- **Configuring, monitoring, and diagnosing protection relationship issues**
After creating and configuring protection relationships, storage administrators can view potential protection reliability issues that are displayed in the Dashboard page, and they 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. Storage administrators can also configure and manage SnapMirror and SnapVault relationships.
- **Performing data restoration**
- **Sending a support bundle to technical support**
Storage administrators can retrieve and send a support bundle to technical support using the maintenance console. Support bundles need to be sent to technical support when the issue requires more detailed diagnosis and troubleshooting than an AutoSupport message provides.

Related concepts

[Monitoring and troubleshooting data availability](#) on page 42

[Creating, monitoring, and troubleshooting protection relationships](#) on page 70

[Prioritizing storage object events using annotations](#) on page 89

Related tasks

[Configuring your environment after deployment](#) on page 26

[Setting up and monitoring an SVM with Infinite Volume without storage classes](#) on page 58

[Managing your Infinite Volume with storage classes and data policies](#) on page 61

[Resolving capacity issues](#) on page 68

[Setting up protection relationships in Unified Manager](#) on page 71

[Restoring data from Snapshot copies](#) on page 87

[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 Unified Manager initial setup.
- You must be logged in as the OnCommand Administrator role to perform this task.

About this task

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

Choices

- [Changing the Unified Manager host name](#) on page 27

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 31

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

- [Adding clusters and viewing the discovery status](#) on page 40

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

Related references

[Unified Manager roles and capabilities](#) on page 148

Changing the Unified Manager host name

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

Before you begin

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

About this task

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

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

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

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

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

Steps

1. [Edit the network settings](#) on page 28

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 29

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 30

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 31

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

About this task

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

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

Steps

1. Click **Administration > Configure Network Settings**.
2. In the **Configure Network Settings** dialog box, modify the host and network settings, as required.
Tip: You can enter multiple comma-separated values in the Secondary DNS Address and Search Domains fields.
3. Click **Update**.

After you finish

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

Related tasks

[Changing the Unified Manager host name](#) on page 27

Generating an HTTPS security certificate

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

Before you begin

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

About this task

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

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > HTTPS**.
3. Click **Regenerate HTTPS Certificate**.

Important: You must restart the Unified Manager virtual machine before the new certificate takes effect. You can use the **System Configuration** option in the NetApp maintenance console.

After you finish

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

If you need to reactivate performance monitoring updates from Performance Manager servers to the Unified Manager server, you must delete the connections that were invalidated by this task and reconfigure new connections.

Related tasks

[Changing the Unified Manager host name](#) on page 27

Viewing the HTTPS security certificate

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

Before you begin

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

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > HTTPS**.
3. Click **View HTTPS Certificate**.

The Subject DN field should display the same host name or fully qualified domain name (FQDN) that is displayed in the Configure Network Settings dialog box. The IP addresses should also be the same in the certificate and in the network settings.

To view more detailed information about the security certificate, you can view the connection certificate in your browser.

Related tasks

[Changing the Unified Manager host name](#) on page 27

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 27

Configuring Unified Manager to send alert notifications

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

Before you begin

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

About this task

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

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

Steps

1. [Configure notification settings](#) on page 32

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 33

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 34

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

4. [Edit global threshold settings](#) on page 35

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 37

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 38

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

Steps

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

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

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

Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 31

Enabling remote authentication

Using Open LDAP or Active Directory, you can enable remote authentication so that the management server can communicate with your authentication servers and so that users of the authentication servers can use Unified Manager to manage the storage objects and data.

Before you begin

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

About this task

If remote authentication is disabled, remote users or groups can no longer access Unified Manager.

The only two supported remote authentication methods are Active Directory and Open LDAP. LDAPS is not supported.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > Authentication**.
3. Select **Enable Remote Authentication**.

If you are using Active Directory Authentication Service, you can enter the Administrator Name using one of the following formats:

- domainname/username
- username@domainname
- Bind Distinguished Name (using appropriate LDAP notation)

4. Optional: Add authentication servers and test the authentication.
5. Click **Save and Close**.

Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 31

[Adding authentication servers](#) on page 34

Adding authentication servers

You can add authentication servers and enable remote authentication on the management server to enable remote users within the authentication server to access Unified Manager.

Before you begin

- The following information must be available:
 - Host name or IP address of the authentication server
 - Port number of the authentication server
- You must have enabled remote authentication and configured your authentication service so that the management server can authenticate remote users or groups in the authentication server.
- You must be logged in as the OnCommand Administrator role to perform this task.

About this task

If the authentication server that you are adding is part of a high-availability pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > Authentication**.
3. In the Servers area, click **Add**.
4. In the **Add Authentication Server** dialog box, specify either the host name or IP address of the server, and the port details.
5. Click **Add**.

Result

The authentication server that you added is displayed in the Servers area.

After you finish

Perform a test authentication to confirm that you are able to authenticate users in the authentication server that you added.

Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 31

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 35
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 36
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 37
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 31

Configuring global aggregate threshold values

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

Before you begin

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

About this task

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

Steps

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

Configuring global volume threshold values

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

Before you begin

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

About this task

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

Steps

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

Editing unmanaged relationship lag threshold settings

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

Before you begin

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

Steps

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

Adding a user

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

Before you begin

- To add a remote user or group, you must have enabled remote authentication and configured your authentication server.
- You must be logged in as the OnCommand Administrator role to perform this task.

About this task

If you add a group from active directory, then all direct members and nested subgroups can authenticate to Unified Manager. If you add a group from OpenLDAP or Other authentication services, then only direct members of that group can authenticate to Unified Manager.

Steps

1. Click **Administration > Manage Users**.
2. In the **Manage Users** page, click **Add**.
3. In the **Add User** dialog box, select the type of user that you want to create and enter the required information.

When entering the required user information, you must specify an email address unique to that user. Specifying email addresses shared by multiple users must be avoided.

4. Click **Add**.

Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 31

Adding an alert

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

Before you begin

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

About this task

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

Steps

1. Click **Administration > Manage Alerts**.
2. In the **Alerts** page, click **Add**.
3. In the **Add Alert** dialog box, perform the following steps:
 - a) Click **Name** and enter a name and description for the alert.
 - b) Click **Resources** and select the resources to be included or excluded from the alert.

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

Note: The text string that you specify is case-sensitive.

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

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

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

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

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

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

You can also choose to notify users through SNMP traps.

4. Click **Save**.

Example for adding an alert

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

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

Perform the following steps in the Add Alert dialog box:

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

7. Click **Save**.

Related concepts

Event state definitions on page 19

Related tasks

Configuring Unified Manager to send alert notifications on page 31

Related references

Description of event severity types on page 101

Description of event impact levels on page 101

Adding clusters

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

Before you begin

- The following information must be available:
 - Host name or cluster management IP address
The host name is the FQDN or short name that Unified Manager uses to connect to the cluster. This host name must resolve to the cluster-management IP address.
The cluster-management IP address must be the cluster-management LIF of the admin Storage Virtual Machine (SVM). If you use a node-management LIF, the operation fails.
 - Data ONTAP administrator user name and password
 - Type of protocol (HTTP or HTTPS) that can be configured on the cluster and the port number of the cluster
- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

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

5. If HTTPS is selected, perform the following steps:
 - a) In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information of the cluster.
 - b) Click **Yes**.

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

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

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

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

Result

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

Changing the local user password

You can change your login password to prevent potential security risks.

Before you begin

You must be logged in as a local user.

About this task

The passwords for the maintenance user and for the remote user cannot be changed from the web UI. To change the maintenance user password, use the Unified Manager maintenance console. To change the remote user password, contact your password administrator.

Steps

1. Log in to Unified Manager.
2. Click `user_name` > **Change Password**.

The **Change Password** option is not displayed if you are a remote user.
3. In the **Change Password** dialog box, enter the details as required.
4. Click **Save**.

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 15

Related tasks

[Resolving a flash card offline condition](#) on page 42

[Scanning for and resolving storage failover interconnect link down conditions](#) on page 44

[Resolving volume offline issues](#) on page 47

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

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 flash card is offline, availability of stored data is impeded because the performance of the cluster node on which it is installed is impaired. You can perform the following steps to localize and identify the potential problem:

Steps

1. From the **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 is displayed.
2. On the **Event details** page, you can review the information displayed in the Cause field and perform one or more of the following tasks:
 - Assign the event to an administrator. [Assigning events](#) on page 96
 - 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 43
 - Acknowledge the event. [Acknowledging and resolving events](#) on page 97

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

About this task

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

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

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

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

Steps

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

In this example, the source of the event is the “alpha-node” cluster node. Clicking that node name displays the HA Details tab on the Nodes tab of the Cluster details page for the affected cluster. The displayed HA Details tab displays information about the 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 a detailed description of the event and often advice to remedy the condition indicated by the event.

After you finish

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

Related references

[Event details page](#) on page 98

[Cluster details page](#) on page 131

[Unified Manager roles and capabilities](#) on page 148

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

About this task

If storage failover interconnections between 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) 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 are displayed 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 is displayed.

- b) On the **Event details** page, you can perform one or more of the following tasks:
 - Review the error message in the Cause field and evaluate the issue. [Performing corrective action for storage failover interconnect links down](#) on page 46
 - Assign the event to an administrator. [Assigning events](#) on page 96
 - Acknowledge the event. [Acknowledging and resolving events](#) on page 97

Related references

[Event details page](#) on page 98

[Unified Manager roles and capabilities](#) on page 148

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

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 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 , aardvark and bonobo, on the Nodes tab of the Cluster details page, and displays other events that recently occurred on the affected .

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 nodes.

Related references

[Event details page](#) on page 98

[Cluster details page](#) on page 131

[Unified Manager roles and capabilities](#) on page 148

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

About this task

Volumes might be reported offline for several reasons:

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

You can use the 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 is displayed.

2. On that page, check the notes for any indication that the SVM administrator has taken the volume in question offline.
3. On the **Event details** page, you can review the information for one or more of the following tasks:
 - Review the information displayed in the Cause field for possible diagnostic guidance. In this example, the information in the Cause field informs you only that the volume is offline.
 - Check the Notes and Updates area for any indication that the SVM administrator has deliberately taken the volume in question offline.
 - Click the source of the event, in this case the volume that is reported offline, to get more information about that volume. [Performing corrective action for volume offline conditions](#) on page 48
 - Assign the event to an administrator. [Assigning events](#) on page 96
 - Acknowledge the event or, if appropriate, mark it as resolved. [Acknowledging and resolving events](#) on page 97

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

About this task

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

Starting at the offline volume's 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 partner has failed also.
See [Determining if a volume offline condition is caused by a down cluster node](#) on page 49.
- Click **Volume details** page links to determine if the volume is offline and its host Storage Virtual Machine (SVM) is stopped because the cluster node hosting the root volume of that SVM is down.
See [Determining if a volume is offline and SVM is stopped because a cluster node is down](#) on page 50.
- 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 52.

Related references

[Unified Manager roles and capabilities](#) on page 148

[Volume details page](#) on page 102

[Storage Virtual Machine details page](#) on page 116

[Cluster details page](#) on page 131

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 partner is unsuccessful.

Before you begin

The following requirements must be met:

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

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 SVM in the **Related Devices** pane of the offline volume's **Volume details** page.

The Storage Virtual Machine details page displays information about the offline volume's hosting Storage Virtual Machine (SVM).

2. In the **Related Devices** pane of the **Storage Virtual Machine 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 SVM.

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

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

4. On the **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 the name of its hosting cluster node.

The Nodes tab on the Cluster details page displays the state of the 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 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 SVM 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 Storage Virtual Machine (SVM) is stopped due to the cluster node hosting the root volume of that SVM being down.

Before you begin

The following requirements must be met:

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

About this task

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

Steps

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

The Storage Virtual Machine details page displays the “running” or the “stopped” status of the hosting SVM. If the SVM status is running, then the volume offline condition is not caused by the cluster node hosting the root volume of that SVM being down.

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

The information for that SVM is shown in a table.

5. On the Storage Virtual Machines page, click  and then select the option **Root Volume**.

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

6. In the Root Volume column, click the name of the root volume to display the **Storage Virtual Machine details** page for that volume.

If the status of the SVM root volume is (Online), then the original volume offline condition is not caused because the cluster node hosting the root volume of that SVM is down.

7. If the status of the SVM root volume is (Offline), then locate and click the hypertext link displayed under Aggregate in the **Related Devices** pane of the SVM root volume's **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 of nodes to which the SVM 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 SVM offline condition, which itself is caused by the cluster node that hosts the root volume of that SVM 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 assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.

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 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 is displayed.

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

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

After you finish

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

Diagnosing performance issues using Unified Manager and Performance Manager

You can use Unified Manager and Performance Manager to diagnose performance issues on resources monitored by the two applications.

Setup tasks for a connection between Performance Manager and Unified Manager

Setting up a connection between a Performance Manager server and the Unified Manager server includes creating a specialized Events Publisher user in the Unified Manager web UI and enabling the Unified Manager server connection in the maintenance console of that Performance Manager server.

Creating a user with Event Publisher role privileges

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

Before you begin

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

About this task

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

Steps

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

Result

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

Related concepts

[Purpose of a connection between Performance Manager and Unified Manager](#) on page 23

[Connections between multiple Performance Manager servers and Unified Manager](#) on page 24

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

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

Before you begin

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

About this task

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

For each connection, complete the following actions:

Steps

1. Log in as the maintenance user to the maintenance console of the Performance Manager server for which you want to create the Unified Manager connection.
2. In the maintenance console, type the number of the menu option labeled "Unified Manager Server Connection" and then type the number of the menu option labeled "Add/Modify Unified Manager Server Connection."

3. When prompted, supply the requested Unified Manager server name or IP address and Unified Manager server port information.

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

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

Result

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

Related concepts

[What performance incidents are](#) on page 20

[Display of performance incidents in Unified Manager](#) on page 20

[Workloads, cluster components, and performance](#) on page 21

Analyzing a performance incident

This workflow provides an example of how you might use Unified Manager connected with Performance Manager to flag and then diagnose performance issues on your managed storage and then use OnCommand System Manager to resolve those issues.

Before you begin

- Unified Manager must be connected with the Performance Manager server that generated the performance incident.
- You must have Storage Administrator or OnCommand Administrator access to Unified Manager.
- You must have Operator, Storage Administrator or OnCommand Administrator access to the Performance Manager server associated with the performance incident that you want to diagnose.
- You must have storage administrator access through OnCommand System Manager to the affected clusters.

About this task

In this example workflow, the performance incident to be analyzed is caused by the following conditions and circumstances:

A Performance Manager server is connected to a Unified Manager server.

The Performance Manager server is monitoring a set of volumes as workloads that are members of a common policy group.

The Unified Manager server is monitoring those same volumes as storage objects.

An increase in the I/O activity for one volume workload causes the policy group to which it belongs to throttle the workload activity for all the members of that policy group.

The throttling causes Performance Manager to generate a performance incident and forward notice of that incident to Unified Manager, which displays that incident in the Performance pane of the Unified Manager Dashboard page.

Steps

1. Your first steps are noting the performance incident in Unified Manager and opening Performance Manager to determine the nature of the incident:
 - a) As storage administrator, you scan the Unified Manager **Dashboard** page to check monitor the general storage health of your managed clusters.

You notice in the Performance pane of the Dashboard page, that a performance incident message is displayed. The message consists of the generic label “performance incident” and the name of the affected cluster component.
 - b) In the Unified Manager **Dashboard** page **Performance** pane, you click the hypertext link of the performance incident that is displayed.

This action opens a separate Performance Manager tab on your browser, with login prompts to the Performance Manager server that discovered the performance incident.
 - c) For follow up use, you note in the URL line of the **Performance Manager** tab, the web address or the IP address of the Performance Manager server.
 - d) In the **Performance Manager** browser tab, you log in to Performance Manager.

Performance Manager displays the Incident Details page for the performance incident in question.
 - e) In the **Summary** section of the **Incident Details** page, you read the Performance Incident **Description**.

In this example workflow, the description reads:
 Volume X is causing Volume X to be slow, due to the policy group limit.

Note: In this example workflow, the same volume workload (Volume X) is displayed as the victim and the bully, because the throttling makes the workload a victim of itself.
 - f) You record the name of the volume for later use.
2. In this example workflow, your next steps are to use the Performance Manager options and tables to determine whether to set a new limit on the policy group, to stop the throttling:
 - a) In the **Workload details** table of the **Incident Details** page, you click the **Activity** column header and select **Peak Deviation in Throttling**.

Your selection causes the volume workloads in the policy group to be sorted by highest deviation of actual activity from their expected activity. The volume workload at the top of

the list has the highest deviation and caused the throttling to occur. The activity is the percentage of the policy group limit used by each workload.

- b) In the **Workloads** column, you click the name of the volume workload at the top of the column.

The Volume Details page for that volume workload is displayed, showing detailed performance data for the selected workload.

- c) In the **Volume Details** page, you select **Break down data by**.

This action displays a selection box for the volume data you want to display.

- d) In the selection box, you select **Response Time** and **Reads/writes/other** (under operation), and then click **Submit**.

The breakdown charts are displayed under the Response Time chart and the Operations chart.

- e) You compare the **Policy Group Impact** chart to the **Response Time** chart to see what percentage of throttling impacted the response time at the time of the incident:

- The policy group has a maximum throughput limit of 1,000 operations per second, which the workloads in it cannot collectively exceed.
- At the time of the incident, the workloads in the policy group had a combined throughput of over 1,200 operations per second, which caused the policy group to throttle their activity back to 1,000 operations per second.
- The Policy Group Impact chart shows that the throttling caused 10% of the total response time, confirming that the throttling caused the incident to occur.

- f) You review the **Cluster components** chart, which displays the breakdown of the total response time by cluster component.

The chart shows the highest response time at the policy group, further confirming that the policy group throttling caused the incident.

- g) You compare the **Reads/writes** chart to the **Reads/writes/other** chart.

Both charts show a high number of read requests with high latency, but the number of requests and amount of latency for write requests is low. These values help you determine whether a high amount of throughput or a high number of operations that increased the response time. You decide to increase the policy group limit on operations.

3. In this example workflow, you use OnCommand System Manager to increase the policy group limit, and then return to Performance Manager to confirm the success of the action:

- a) You use OnCommand System Manager to increase the current limit on the policy group to 1,300 operations per second.
- b) You wait for a spike in storage activity equal to the spike that triggered the original performance incident.
- c) Using your browser and the web address or IP address that you recorded in Step 1c, you return to Performance Manager.
- d) You search for the name of the volume workload that you recorded in Step 1e.

The Volume Details page is displayed.

- e) You select **Break down data by > Operations**.

- f) You click **Submit**.
The Cluster components chart is displayed.
- g) At the bottom of the page, you move your cursor to the change event icon for the policy group limit change.
- h) You compare the **Reads/writes/other** chart to the **Response Time** chart.
The read and write requests are the same, but the throttling has stopped and the response time has decreased.

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

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

Before you begin

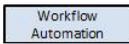
The following requirements must be met:

- 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 assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- 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 SVMs using Unified Manager.
- While performing this task, you are required to switch between two applications: OnCommand Workflow Automation (WFA) and OnCommand Unified Manager.
- The task provides high-level steps.
For details about performing the WFA tasks, see the *OnCommand Workflow Automation* documentation.

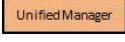
Steps

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

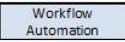
You can enable storage efficiency technologies, such as deduplication and compression, while creating the Infinite Volume.

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

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

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

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

4.  Configure notification alerts and traps to address any availability and capacity issues related to the Infinite Volume.
5. Optional:  Create a disaster recovery (DR) SVM with Infinite Volume, and then configure data protection (DP) by performing the following steps:
 - a) Create a data protection (DP) Infinite Volume by using the appropriate workflow.
 - b) Set up a DP mirror relationship between the source and destination by using the appropriate workflow.

Related tasks

[Adding clusters](#) on page 40

[Editing the Infinite Volume threshold settings](#) on page 60

[Adding an alert](#) on page 38

Related references

[Unified Manager roles and capabilities](#) on page 148

Adding clusters

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

Before you begin

- The following information must be available:
 - Host name or cluster management IP address

The host name is the FQDN or short name that Unified Manager uses to connect to the cluster. This host name must resolve to the cluster-management IP address.

The cluster-management IP address must be the cluster-management LIF of the admin Storage Virtual Machine (SVM). If you use a node-management LIF, the operation fails.
 - Data ONTAP administrator user name and password

- Type of protocol (HTTP or HTTPS) that can be configured on the cluster and the port number of the cluster
- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

1. Click **Storage > Clusters**.
2. From the **Clusters** page, click **Add**.
3. In the **Add Cluster** dialog box, specify the values, such as the host name or IP address of the cluster, user name, password, protocol for communication, and port number.
By default, the HTTPS protocol is selected.
4. Click **Add**.
5. If HTTPS is selected, perform the following steps:
 - a) In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information of the cluster.
 - b) Click **Yes**.

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

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

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

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

Result

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

Editing the Infinite Volume threshold settings

When you need to address any issues in your Infinite Volume's storage space, you can edit the threshold settings of the Infinite Volume's capacity based on your organization's requirements. When

a threshold is crossed, events are generated, and you receive notifications if you have configured alerts for such events.

Before you begin

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

Steps

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

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.

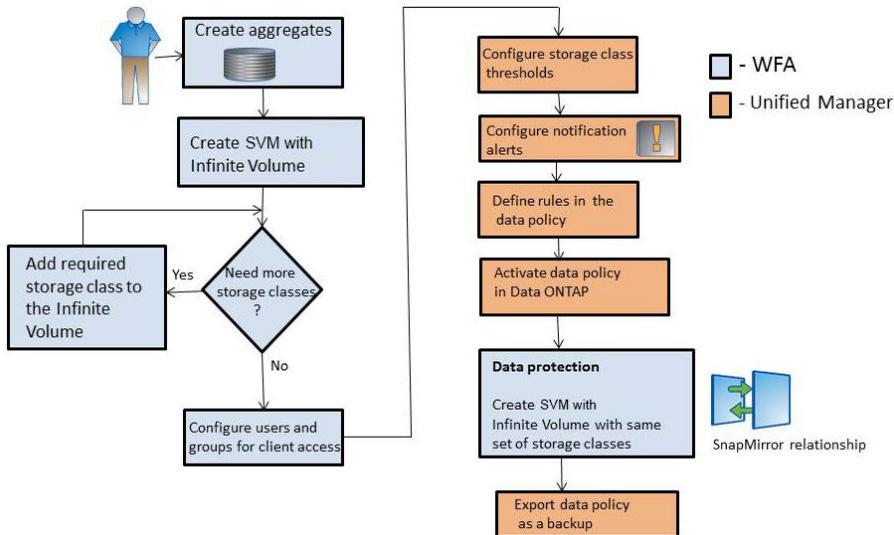
Note: Unified Manager and WFA are separate installations.

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

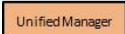
While performing this task, you are required to switch between two applications: OnCommand Workflow Automation (WFA) and OnCommand Unified Manager.

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



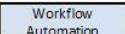
Steps

1. Workflow Automation Customize the predefined workflow to define the required storage classes.
2. Workflow Automation Create an SVM with Infinite Volume with the required number of storage classes by using the appropriate workflow.
3. Unified Manager Add the cluster containing the SVM with Infinite Volume to the Unified Manager database.
 You can add the cluster by providing the IP address or the FQDN of the cluster.
4. Unified Manager *Based on your organization's requirements, modify the thresholds for each storage class* on page 63.
 You should use the default storage class threshold settings to effectively monitor storage class space.
5. Unified Manager *Configure notification alerts and traps to address any availability and capacity issues related to the Infinite Volume* on page 38.

6.  *Set up rules in the data policy, and then activate all the changes made to the data policy* on page 66

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) SVM with Infinite Volume, and then configure a data protection (DP) by performing the following steps:
- Create a data protection (DP) Infinite Volume by using the appropriate workflow.
 - 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 assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.

Steps

- Click **Storage > Storage Virtual Machines**.
- In the Storage Virtual Machines page, select the required SVM with Infinite Volume.
- In the **Storage Virtual Machine details** page, click **Actions > Edit Thresholds**.
- In the **Edit Storage Class Thresholds** dialog box, modify the thresholds as required.
- 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 notification settings such as the email address, SMTP server, and SNMP trap host so that the Unified Manager server can use these settings to send notifications to users when an event is generated.

- The following information must be available: resources and events for which you want to trigger the alert, and user names or email addresses of users you want to notify.
- You must be logged in as the OnCommand Administrator role to perform this task..

About this task

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

Steps

1. Click **Administration > Manage Alerts**.
2. In the **Alerts** page, click **Add**.
3. In the **Add Alert** dialog box, perform the following steps:

- a) Click **Name** and enter a name and description for the alert.
- b) Click **Resources** and select the resources to be included or excluded from the alert.

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

Note: The text string that you specify is case-sensitive.

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

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

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

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

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

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

You can also choose to notify users through SNMP traps.

4. Click **Save**.

Example for adding an alert

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

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

Perform the following steps in the Add Alert dialog box:

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

Related concepts

[Event state definitions](#) on page 19

Related tasks

[Configuring Unified Manager to send alert notifications](#) on page 31

Related references

[Description of event severity types](#) on page 101

[Description of event impact levels](#) on page 101

Creating rules

You can add new rules to your data policy to determine the placement of data that is written to the Infinite Volume. You can create rules either by using rule templates that are defined in Unified Manager or create custom rules.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

Choices

- [Creating rules using templates](#) on page 66
- [Creating custom rules](#) on page 67

Creating rules using templates

You can add new rules by using rule templates defined by Unified Manager to determine the placement of data that is written to the Storage Virtual Machine (SVM) with Infinite Volume. You can create rules based on file types, directory paths, or owners.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click the **Data Policy** tab.
The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Click **Create**.
5. In the **Create Rule** dialog box, choose an appropriate rule template from the drop-down list.
The template is based on three categories: file type, owner, or directory path.

6. Based on the template selected, add necessary conditions in the **Matching Criteria** area.
7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.
8. Click **Create**.
The new rule you created is displayed in the Data Policy tab.
9. Optional: Preview any other changes made to the data policy.
10. Click **Activate** to activate the changes in the rule properties in the SVM.

Creating custom rules

Based on your data center requirements, you can create custom rules and add them to a data policy to determine the placement of data that is written to the SVM with Infinite Volume. You can create custom rules from the Create Rule dialog box without using any existing template.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.
4. Click **Create**.
5. In the **Create Rule** dialog box, select **Custom rule** from the **Template** list.
6. In the **Matching Criteria** area, add conditions as required.
Conditions enable you to create a rule based on file types, directory paths, or owners. A combination of these conditions are the condition sets. For example, you can have a rule: “Place all .mp3 owned by John in bronze storage class.”
7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.
8. Click **Create**.
The newly created rule is displayed in the Data Policy tab.

9. Optional: Preview any other changes made to the data policy.
10. Click **Activate** to activate the changes in the rule properties in the SVM.

Exporting a data policy configuration

You can export a data policy configuration from Unified Manager to a file. For example, after you have taken the required backup, and in the event of a disaster, you can export the data policy configuration from the primary.

Before you begin

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

About this task

The Data Policy tab, which is used while performing this task, is displayed only for SVMs with Infinite Volume.

Steps

1. Click **Storage > Storage Virtual Machines**.
2. In the Storage Virtual Machines page, select the appropriate SVM.
3. Click **Data Policy**.
The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.
4. Click **Export**.
5. In the browser-specific dialog box, specify the location to which the data policy configuration has to be exported.

Result

The data policy configuration is exported as a JSON file in the specified location.

Resolving capacity issues

This workflow provides an example of how you can resolve a capacity issue. In this scenario, you are an administrator or operator and you access the Unified Manager 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 assigned one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator..

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 SVM 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 is displayed.
2. From the **Event details** page, you can perform one or more of the following tasks:
 - Review the error message in the Cause field and click the suggestions under Suggested Remedial Actions to review descriptions of possible remediations. *Performing suggested remedial actions for a full volume* on page 69
 - Click the object name, in this case a volume, in the Source field to get details about the object. *Volume details page* on page 102
 - Look for notes that might have been added about this event. *Adding and reviewing notes associated with an event* on page 96
 - Add a note to the event. *Adding and reviewing notes associated with an event* on page 96
 - Assign the event to another user. *Assigning events* on page 96
 - Acknowledge the event. *Acknowledging and resolving events* on page 97
 - Mark the event as resolved. *Acknowledging and resolving events* on page 97

Related references

Event details page on page 98

Performing suggested remedial actions for a full volume

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

Before you begin

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 98

[Volume details page](#) on page 102

[Aggregate details page](#) on page 140

Creating, monitoring, and troubleshooting protection relationships

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

Related concepts

[How protection relationships are created and protection jobs run](#) on page 17

Related tasks

[Resolving a protection job failure](#) on page 81

[Resolving lag issues](#) on page 85

[Setting up protection relationships in Unified Manager](#) on page 71

[Performing a protection relationship failover and failback](#) on page 76

Setting up protection relationships in Unified Manager

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

Before you begin

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

Steps

1. [Set up OnCommand Workflow Automation](#) on page 71.

OnCommand Workflow Automation must be integrated with Unified Manager before you can set up protection relationships.

2. Depending on the type of protection relationship you want to create, do one of the following:
 - [Create a SnapMirror protection relationship](#) on page 72.
 - [Create a SnapVault protection relationship](#) on page 73.
3. If you want to create a policy for the relationship, depending on the relationship type you are creating, do one of the following:
 - [Create a SnapVault policy](#) on page 74.
 - [Create a SnapMirror policy](#) on page 75.
4. [Create a SnapMirror or SnapVault schedule](#) on page 76.

Pairing OnCommand Workflow Automation with Unified Manager

You can integrate Workflow Automation with Unified Manager over a secure connection. This enables protection features such as SnapMirror and SnapVault configuration workflows, as well as commands for managing SnapMirror relationships.

Before you begin

The following information must be available:

- Name of a database user in Unified Manager
- Host address, port number 443, user name, and password for the OnCommand Workflow Automation setup

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

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

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Add-ons > Workflow Automation**.
3. In the **Set Up OnCommand Workflow Automation** dialog box, type the host name or IP address and the user name and password.
You must use Unified Manager server port 443.
4. Click **Save and Close**.
5. If you use a self-signed certificate, click **Yes** to authorize the security certificate.
The Workflow Automation Options Changed dialog box displays.
6. Click **Yes** to reload the web UI and add the Workflow Automation features.

Related tasks

[Setting up protection relationships in Unified Manager](#) on page 71

[Performing a protection relationship failover and failback](#) on page 76

Creating a SnapMirror protection relationship from the Volume details page

You can create a SnapMirror relationship using the Volume details page so that data replication is enabled for protection purposes. SnapMirror replication enables you to restore data from the destination volume in the event of data loss on the source for volumes running Data ONTAP 8.2 or later.

Before you begin

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

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

Steps

1. In the **Protection** tab of the **Volume details** page, right-click in the topology view the name of a volume that you want to protect.
2. Select **Protect > SnapMirror** from the menu.
The Configure Protection dialog box is displayed.
3. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
4. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.
You are returned to the Volume details page.
7. Click the protection configuration job link at the top of the **Volume details** page.
The jobs tasks and details are displayed in the Job details page.
8. In the **Job details** page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
9. When the job tasks are complete, click **Back** on your browser to return to the **Volume details** page.
The new relationship is displayed in the Volume details page topology view.

Related tasks

[Setting up protection relationships in Unified Manager](#) on page 71

[Pairing OnCommand Workflow Automation with Unified Manager](#) on page 71

Creating a SnapVault protection relationship from the Volume details page

You can create a SnapVault relationship using the Volume details page so that data backups are enabled for protection purposes on volumes running Data ONTAP 8.2 or later.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges

- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Protection** tab of the **Volume details** page, right-click a volume in the topology view that you want to protect.
2. Select **Protect > SnapVault** from the menu.
The Configure Protection dialog box is launched.
3. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary resource information.
4. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.
You are returned to the Volume details page.
7. Click the protection configuration job link at the top of the **Volume details** page.
The Job details page is displayed.
8. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
When the job tasks are complete, the new relationships are displayed in the Volume details page topology view.

Related tasks

[Setting up protection relationships in Unified Manager](#) on page 71

[Pairing OnCommand Workflow Automation with Unified Manager](#) on page 71

Creating SnapVault policies

You can create a new SnapVault policy to set the priority for a SnapVault transfer. You use policies to maximize the efficiency of transfers from the primary to the secondary in a protection relationship.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to enable this operation.
- You must have already completed Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

The SnapVault tab is displayed.

2. In the **Policy Name** field, type the name that you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority that you want to assign to the policy.
4. Optional: In the **Comment** field, enter a comment for the policy.
5. In the **Replication Label** area, add or edit a replication label, as necessary.
6. Click **Create**.

The new policy is displayed in the Create Policy drop-down list.

Related tasks

[Setting up protection relationships in Unified Manager](#) on page 71

[Pairing OnCommand Workflow Automation with Unified Manager](#) on page 71

Creating SnapMirror policies

You can create a SnapMirror policy to specify the SnapMirror transfer priority for protection relationships. SnapMirror policies enable you to maximize transfer efficiency from the source to the destination by assigning priorities so that lower-priority transfers are scheduled to run after normal-priority transfers.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to enable this operation.
- This task assumes that you have already completed the Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapMirror** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

The Create SnapMirror Policy dialog box is displayed.

2. In the **Policy Name** field, type a name you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority you want to assign to the policy.
4. In the **Comment** field, enter an optional comment for the policy.
5. Click **Create**.

The new policy is displayed in the SnapMirror Policy drop-down list.

Related tasks

[Setting up protection relationships in Unified Manager](#) on page 71

[Pairing OnCommand Workflow Automation with Unified Manager](#) on page 71

Creating SnapMirror and SnapVault schedules

You can create basic or advanced SnapMirror and SnapVault schedules to enable automatic data protection transfers on a source or primary volume so that transfers take place more frequently or less frequently, depending on how often the data changes on your volumes.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task..
- You must have already completed the Destination Information area in the Configure Protection dialog box.
- You must have set up Workflow Automation to perform this task.

Steps

1. From the **SnapMirror** tab or **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Schedule** link in the **Relationship Settings** area.
The Create Schedule dialog box is displayed.
2. In the **Schedule Name** field, type the name you want to give to the schedule.
3. Select one of the following:
 - **Basic**
Select if you want to create a basic interval-style schedule.
 - **Advanced**
Select if you want to create a cron-style schedule.
4. Click **Create**.
The new schedule is displayed in the SnapMirror Schedule or SnapVault Schedule drop-down list.

Related tasks

[Setting up protection relationships in Unified Manager](#) on page 71

[Pairing OnCommand Workflow Automation with Unified Manager](#) on page 71

Performing a protection relationship failover and failback

When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection

destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Before you begin

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

Steps

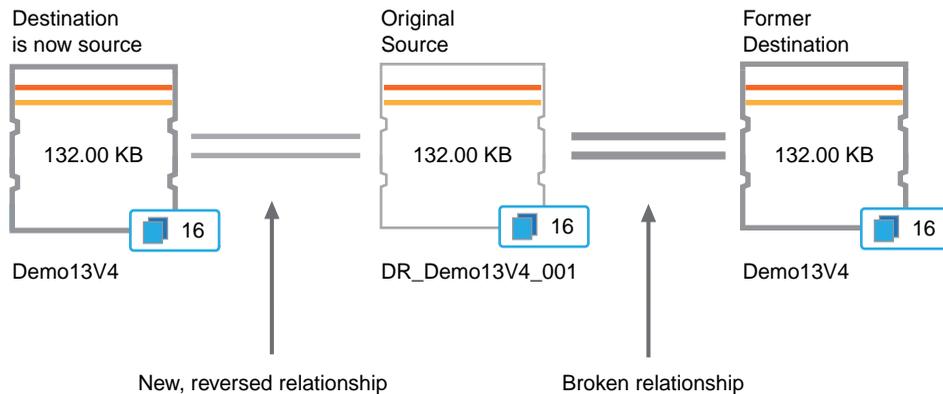
1. [Break the SnapMirror relationship](#) on page 78.

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

2. [Reverse the protection relationship](#) on page 78.

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

The reversed relationship looks similar to a cascaded relationship:



3. [Break the reversed SnapMirror relationship](#) on page 78.

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

4. [Remove the relationship](#) on page 80.

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

5. [Resynchronize the relationship](#) on page 80.

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

Breaking a SnapMirror relationship from the Volume details page

You can break a protection relationship from the Volume details page and stop data transfers between a source and destination volume in a SnapMirror relationship. You might break a relationship when you want to migrate data, for disaster recovery, or for application testing. The destination volume is changed to a read-write volume. You cannot break a SnapVault relationship.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

Steps

1. In the **Protection** tab of the **Volume details** page, select from the topology the SnapMirror relationship you want to break.
2. Right-click the destination and select **Break** from the menu.
The Break Relationship dialog box is displayed.
3. Click **Continue** to break the relationship.
4. In the topology, verify that the relationship is broken.

Related tasks

[Performing a protection relationship failover and failback](#) on page 76

[Pairing OnCommand Workflow Automation with Unified Manager](#) on page 71

Reversing protection relationships from the Volume details page

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- Your system must be running Data ONTAP 8.2 or later.
- Workflow Automation must be set up to perform this operation.

- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.
- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on the reverse resynchronization relationship are the same as those on the original protection relationship.
If policies and schedules do not exist, they are created.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the SnapMirror relationship on which you want to reverse the source and destination, and right-click it.
2. Select **Reverse Resync** from the menu.
The Reverse Resync dialog box is displayed.
3. Verify that the relationship displayed in the **Reverse Resync** dialog box is the one for which you want to perform the reverse resynchronization operation, and then click **Submit**.
The Reverse Resync dialog box is closed and a job link is displayed at the top of the Volume details page.
4. Optional: Click **View Jobs** on the **Volume details** page to track the status of each reverse resynchronization job.
A filtered list of jobs is displayed.
5. Optional: Click the Back arrow on your browser to return to the **Volume details** page.
The reverse resynchronization operation is finished when all job tasks are completed successfully.

Related tasks

[Performing a protection relationship failover and failback](#) on page 76

[Pairing OnCommand Workflow Automation with Unified Manager](#) on page 71

Removing a protection relationship from the Volume details page

You can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must have set up Workflow Automation to perform this task.

Steps

1. In the **Protection** tab of the **Volume details** page, select from the topology the SnapMirror relationship you want to remove.
2. Right-click the name of the destination and select **Remove** from the menu.
The Remove Relationship dialog box is displayed.
3. Click **Continue** to remove the relationship.
The relationship is removed from the Volume details page.

Resynchronizing protection relationships from the Volume details page

You can resynchronize data on a SnapMirror or SnapVault relationship that was broken and then the destination was made read/write so that data on the source matches the data on the destination. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Before you begin

- You must be assigned either the OnCommand Administrator role or the Storage Administrator role to perform this task.
- You must be running Data ONTAP 8.2 or later.
- You must have set up OnCommand Workflow Automation to perform this operation.

Steps

1. From the **Protection** tab of the **Volume details** page, locate in the topology the protection relationship that you want to resynchronize and right-click it.
2. Select **Resynchronize** from the menu.

Alternatively, from the **Actions** menu, select **Relationship > Resynchronize** to resynchronize the relationship for which you are currently viewing the details.

The Resynchronize dialog box is displayed.

3. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
4. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.
The Select Source Snapshot Copy dialog box is displayed.
5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
6. Click **Submit**.
You are returned to the Resynchronize dialog box.
7. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
8. Click **Submit** to begin the resynchronization job.
The resynchronization job is started, you are returned to the Volume details page and a jobs link is displayed at the top of the page.
9. Optional: Click **View Jobs** on the **Volume details** page to track the status of each resynchronization job.
A filtered list of jobs is displayed.
10. Optional: Click the Back arrow on your browser to return to the **Volume details** page.
The resynchronization job is finished when all job tasks successfully complete.

Related tasks

[Performing a protection relationship failover and failback](#) on page 76

[Pairing OnCommand Workflow Automation with Unified Manager](#) on page 71

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 148.

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_svm:/cluster2_src_vol2 - Protection Job Failed`.

The Event details page for the failed protection job displays.

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

See [Identifying the problem and performing corrective actions for a failed protection job](#) on page 82.

Related references

[Unified Manager roles and capabilities](#) on page 148

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 this task.

About this task

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

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

This message provides the following information:

- A backup or mirror job did not complete successfully.
The job involved a protection relationship between the source volume `cluster2_src_vol2` on the virtual server `cluster2_src_svm` and the destination volume `managed_svc2_vol3` on the virtual server named `cluster3_dst_svm`.

- A Snapshot copy job failed for 0426cluster2_src_vol2snap on the source volume cluster2_src_svm:/cluster2_src_vol2.

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

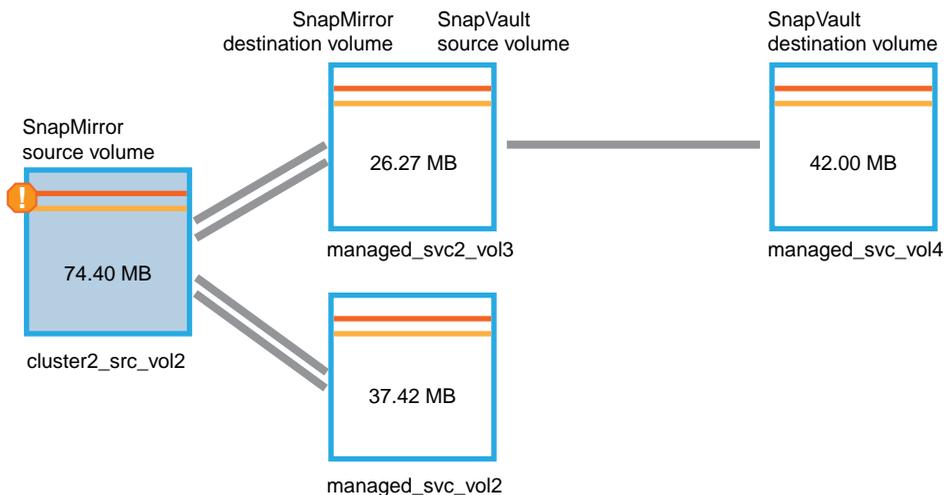
2. You decide that you want to try to resolve the event, so you do the following:
 - a) Click the **Assign To** button and select **Me** from the menu.
 - b) Click the **Acknowledge** button so that you do not continue to receive repeat alert notifications, if alerts were set for the event.
 - c) Optionally, 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 volume autogrow has been disabled and that a volume space guarantee has been set.

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

9. You scroll down to the **Events** list and see that Protection Job Failed, Volume Days Until Full, and Volume Space Full events were generated.

10. In the **Events** list, you click the **Volume Space Full** event to get more information, having decided that this event seems most relevant to your capacity issue.

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

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

12. Below the **Summary** area, you see Suggested Corrective Actions.

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

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

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

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

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

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

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

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

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

14. Go to System Manager or access the 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 96

[Assigning events](#) on page 96

[Acknowledging and resolving events](#) on page 97

Related references

[Job details page](#) on page 146

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 they exist, to find solutions.

Before you begin

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

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

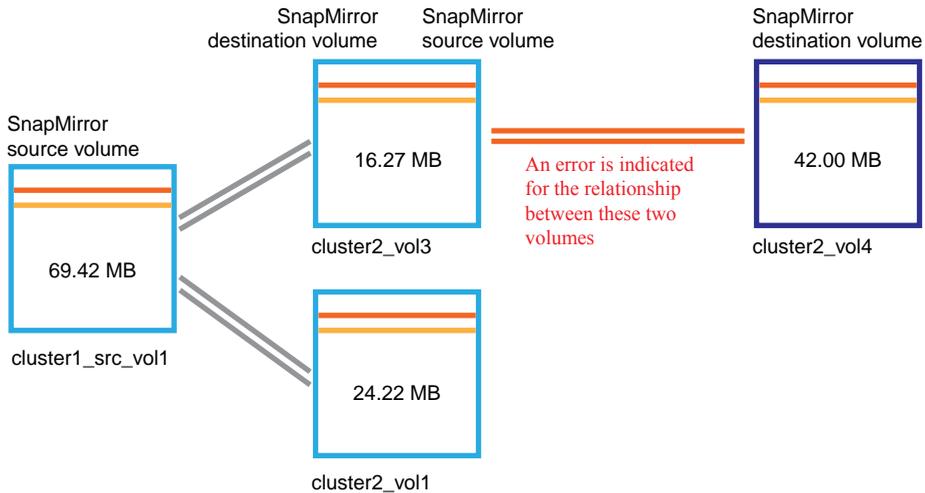
1. In the **Protection** pane on the **Dashboard** page, locate the SnapMirror relationship lag error and click it.
The Event details page for the lag error event is displayed.
2. From the **Event details** page you can perform one or more of the following tasks:
 - Review the error message in the Cause 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.

- 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.

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

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



- 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.

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

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

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

Related tasks

[Adding and reviewing notes about an event](#) on page 96

[Assigning events](#) on page 96

[Acknowledging and resolving events](#) on page 97

Related references

[Event details page](#) on page 98

[Unified Manager roles and capabilities](#) on page 148

Restoring data from Snapshot copies

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

About this task

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

Step

1. Restore data using one of the following tasks:
 - [Restore data from the Volume details page](#) on page 87.
 - [Restore data from the Volumes page](#) on page 88.

Restoring data using the Volume details page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Volume details page. Restoring data from volumes that have a version of Data ONTAP that is earlier than 8.2 is not supported.

Before you begin

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

About this task

You cannot restore NTFS file streams.

The restore option is not available in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Protection** tab of the **Volume details** page, right-click in the topology view the name of the volume that you want to restore.
2. Select **Restore** from the menu.

Alternatively, select **Restore** from the **Actions** menu to protect the current volume for which you are viewing the details.

The Restore dialog box is displayed.

3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.

4. Select the items you want to restore.

You can restore the entire volume, or you can specify folders and files you want to restore.

5. Select the location to which you want the selected items restored: either **Original Location** or **Alternate Existing Location**.

6. If you select an alternate existing location, do one of the following:

- In the Restore Path text field, type the path of the location to which you want to restore the data and then click **Select Directory**.
- Click **Browse** to launch the Browse Directories dialog box and complete the following steps:
 - a. Select the cluster, SVM, and volume to which you want to restore.
 - b. In the Name table, select a directory name.
 - c. Click **Select Directory**.

7. Click **Restore**.

The restore process begins.

Restoring data using the Volumes page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Volumes page. Restoring data from volumes that have a version of Data ONTAP that is earlier than 8.2 is not supported.

Before you begin

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

About this task

You cannot restore NTFS file streams.

The restore option is not available in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Volumes** page, select a volume from which you want to restore data.
2. From the toolbar, click **Restore**.
The Restore dialog box is displayed.
3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
4. Select the items you want to restore.
You can restore the entire volume, or you can specify folders and files you want to restore.
5. Select the location to which you want the selected items restored; either **Original Location** or **Alternate Location**.
6. Click **Restore**.
The restore process begins.

Prioritizing storage object events using annotations

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

Related concepts

[Understanding more about annotations](#) on page 91

[What annotations are](#) on page 19

Related tasks

[Creating rules to annotate storage objects](#) on page 89

[Viewing annotations](#) on page 90

[Removing annotations for storage objects](#) on page 90

Creating rules to annotate storage objects

You can create rules to annotate clusters, volumes, and Storage Virtual Machines (SVMs) using annotations. Annotating storage objects enables you to view and manage object-related events.

Before you begin

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

Steps

1. Click **Administration > Manage Annotations**.
2. In the **Auto-annotation Rules** area, click **Configure**.
3. In the **Configure Workflow Annotation Rules** dialog box, enter the rule based on your criteria.
4. Click **Validate** to validate the syntax of the rule.

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

5. Click **Save**.
6. Verify that the storage objects you tagged are displayed in the **Members** tab of the **Manage Annotations** page.

Viewing annotations

You can view a list of the clusters, volumes, and Storage Virtual Machines (SVMs) that are annotated so that you can manage these objects based on their priority in your environment.

Before you begin

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

Steps

1. Click **Administration > Manage Annotations**.
2. View the annotation for your storage objects by selecting **Mission Critical**, **High**, or **Low**.

Removing annotations for storage objects

You can remove existing annotations for storage objects when you want to change or remove the prioritization of the objects.

Before you begin

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

Steps

1. Click **Administration > Manage Annotations**.
2. Delete the rule for configuring annotations.
3. Click **Save**.

Understanding more about annotations

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

Description of annotation types

Annotations are categorized as three types, mission critical, high, and low. These categories are the only types that enable you to annotate storage objects based on the priority of data that they contain. Annotation types are non-editable.

Mission critical	This annotation is applied to storage objects that contain mission-critical data. For example, objects that contain production applications can be considered as mission critical.
High	This annotation is applied to storage objects that contain high priority data. For example, objects that are hosting business applications can be considered high priority.
Low	This annotation is applied to storage objects that contain low priority data. For example, objects that are on secondary storage, such as backup and mirror destinations, might be of low priority.

Related concepts

[Understanding more about annotations](#) on page 91

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 logged in as the Maintenance User to perform this task.

About this task

For more information about the maintenance console and support bundles, see [Using the maintenance console](#) on page 150.

Unified Manager stores two generated support bundles at one time.

Steps

1. [Accessing the maintenance console using Secure Shell](#) on page 92
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 93
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 93
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 94
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 95
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 be logged in as the Maintenance User to perform this task.

Steps

1. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance.
2. Log in to the maintenance console using your maintenance user name and password.
After 15 minutes of inactivity, the maintenance console logs you out.

Related tasks

[Using the maintenance console](#) on page 150

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

About this task

Unified Manager stores two generated support bundles at one time.

Steps

1. In the maintenance console **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 151

Related references

[Unified Manager roles and capabilities](#) on page 148

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 for the support bundle you want to copy.
5. Select the support bundle you want to copy and use the tool to copy the file from the Unified Manager server to your local system.

Related information

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

WinSCP - <http://winscp.net>

Retrieving the support bundle using a UNIX or Linux client

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

Before you begin

You must be the maintenance user to perform this task.

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

Steps

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

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

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

SFTP	<code>sftp <maintenance-user>@<vApp-name-or-ip>:/support/ support_bundle_file_name.7z <destination-directory></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 requires more detailed diagnosis and troubleshooting than an AutoSupport message provides.

Before you begin

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

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

Steps

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

Related information

NetApp Support Site: support.netapp.com

Tasks and information related to several workflows

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

Adding and reviewing notes about an event

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

Before you begin

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

Steps

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

Assigning events

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

Before you begin

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

Steps

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

If you want to assign the event to...	Then do this...
Yourself	Click Assign To > Me .
Another user	<ol style="list-style-type: none"> <li data-bbox="417 345 780 369">a. Click Assign To > Another user. <li data-bbox="417 392 1208 444">b. In the Assign Owner dialog box, enter the user name, or select a user from the drop-down list. <li data-bbox="417 466 841 520">c. Click Assign. An email notification is sent to the user. <p data-bbox="482 543 1208 591">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.</p>

Acknowledging and resolving events

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

Before you begin

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

About this task

You can acknowledge and resolve multiple events simultaneously.

Steps

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

If you want to...	Do this...
Acknowledge and mark a single event as resolved	<ol style="list-style-type: none"> <li data-bbox="536 1329 780 1354">a. Click the event name. <li data-bbox="536 1376 1173 1400">b. From the Event details page, determine the cause of the event. <li data-bbox="536 1423 774 1447">c. Click Acknowledge. <li data-bbox="536 1470 908 1494">d. Take appropriate corrective action. <li data-bbox="536 1517 827 1541">e. Click Mark As Resolved.

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

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

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

Event details page

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

- [Command buttons](#) on page 98
- [Summary area](#) on page 99
- [Notes and Updates area](#) on page 100
- [Possible Causes](#) on page 100
- [Resources that Might be Impacted](#) on page 100
- [Suggested Corrective Actions area](#) on page 100

Command buttons

The command buttons enable you to perform the following tasks:

Assign To	Me	Assigns the event to you.
	Another user	Opens the Assign Owner dialog box, which enables you to assign or reassign the event to other users.
		When you assign an event to a user, the user's name and the time when the event was assigned are added in the events list for the selected events.
		Note: You can also unassign events by leaving the ownership field blank.

Acknowledge	Acknowledges the selected events so that you do not continue to receive repeat alert notifications.
Mark As Resolved	Enables you to change the event state to resolved.
Add Alert	Displays the Add Alert dialog box, which enables you to add an alert for the selected event.
View Events	Navigates to the Events page.

Summary area

You can view the following event details:

Severity	Displays the severity of the event. The event severity types are Critical () , Error () , Warning () , and Information () .
State	Displays the event state: New, Acknowledged, Resolved, or Obsolete.
Impact Level	Displays whether the event is categorized as an incident, risk, or an informational event.
Impact Area	Displays whether the event is a capacity, availability, protection, performance, or configuration related event.
Obsoleted Cause	Displays the reason the event is now obsolete.
Source	Displays the full name of the object along with the type of object with which the event is associated. The value is displayed as “Unknown” when Data ONTAP does not provide a valid user name because of SecD errors.
Source Type	Displays the object type (for example, Storage Virtual Machine (SVM), volume, or qtrees) with which the event is associated.
Acknowledged By	Displays the name of the person who acknowledged the event and the time when the event was acknowledged. This field is blank if the event is not acknowledged.
Resolved By	Displays the name of the person who resolved the event and the time when the event was resolved. This field is blank if the event is not resolved.
Assigned To	Displays the name of the person who is assigned to work on the event.
Triggered Time	Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp when the event was generated is displayed.
Trigger Condition	Displays information about the cause of the event.
Alert Settings	The following information about alerts is displayed:

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

Alerts that are disabled are not displayed.

Notes and Updates area

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

Post Enables you to display the information that you added.

Possible Causes area

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

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

Resources that Might be Impacted area

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

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

Suggested Corrective Action area

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

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

Related tasks

[Performing diagnostic actions for volume offline conditions](#) on page 48

Performing suggested remedial actions for a full volume on page 69

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

- Availability** Availability events notify you if a storage object goes offline, if a protocol service goes down, if an issue with storage failover occurs, or if an issue with hardware occurs.
- Capacity** Capacity events notify you if your aggregates, volumes, or LUNs are approaching or have reached a size threshold, or if the rate of growth is unusual for your environment.

- Configuration** Configuration events inform you of the discovery, deletion, addition, removal, or renaming of your storage objects. Configuration events have an impact level of Event and a severity type of Information.
- Performance** Performance events, also called incidents, notify you of resource, configuration, or activity conditions on your storage cluster that might adversely affect the speed of data storage input or retrieval on your monitored SVM and volumes. Description of performance impact events is provided in OnCommand Performance Manager help.
- Protection** Protection events notify you of incidents or risks involving SnapMirror relationships, issues with destination capacity, problems with SnapVault relationships, or issues with protection jobs. Any Data ONTAP object (especially aggregates, volumes, and Storage Virtual Machines (SVMs)) that host secondary volumes and protection relationships are categorized in the protection impact area.

Volume details page

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

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

- [Command buttons](#) on page 102
- [Capacity tab](#) on page 104
- [Efficiency tab](#) on page 107
- [Configuration tab](#) on page 108
- [Protection tab](#) on page 109
- [History area](#) on page 114
- [Events list](#) on page 115
- [Related Devices pane](#) on page 115
- [Related Alerts pane](#) on page 116
- [Annotations pane](#) on page 116

Command buttons

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

- Actions**
 - Add Alert
Enables you to add an alert to the selected volume.
 - Edit Thresholds
Enables you to modify the threshold settings for the selected volume.

- **Protect**
Enables you to create either SnapMirror or SnapVault relationships for the selected volume.
- **Relationship**
Enables you to execute the following protection relationship operations:
 - **Edit**
Launches the Edit Relationship dialog box which enables you to change existing SnapMirror policies, schedules, and maximum transfer rates for an existing protection relationship.
 - **Abort**
Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer.
 - **Quiesce**
Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress will complete before the relationship is quiesced.
 - **Break**
Breaks the relationship between the source and destination volume and changes the destination to a read-write volume.
 - **Remove**
Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed. This operation cannot be undone.
 - **Resume**
Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used if one exists.
 - **Resynchronize**
Enables you to resynchronize a previously broken relationship.
 - **Initialize/Update**
Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform a manual update if the relationship is already initialized.
 - **Reverse Resync**
Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by the contents on the destination, and any data that is newer than the data on the common Snapshot copy is deleted.
- **Restore**
Enables you to restore data from one volume to another volume.

View Volumes Enables you to navigate to the Volumes page.

Capacity tab

The Capacity tab displays details about the selected volume, such as its capacity, threshold settings, quota capacity, and information about any volume move operation:

Capacity Displays the display capacity details of the volume.

- **Snapshot Overflow**
Displays the data space that is consumed by the Snapshot copies.
- **Used**
Displays the space used by data in the volume.
- **Warning**
Indicates that the space in the volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.
- **Error**
Indicates that the space in the volume is full. If this threshold is breached, the Space Full event is generated.
- **Unusable**
Indicates that the Thin-Provisioned Volume Space At Risk event is generated and the space in the thinly-provisioned volume is at risk because of aggregate capacity issues. The unusable capacity is displayed only for thinly-provisioned volumes.
- **Data graph**
Displays the total data capacity and the used data capacity of the volume.
If autogrow is enabled, the data graph also considers the space available in the aggregate. The data graph displays the effective storage space that can be used by data in the volume, which can be one of the following:
 - Actual data capacity of the volume for the following conditions:
 - Autogrow is disabled
 - Autogrow-enabled volume has reached the maximum size
 - Autogrow-enabled thickly provisioned volume cannot grow further
 - Data capacity of the volume after considering the maximum volume size (for thinly provisioned volumes and for thickly provisioned volumes when the aggregate has space for the volume to reach maximum size)
 - Data capacity of the volume after considering the next possible autogrow increment (for thickly provisioned volumes that can have at least one autogrow increment)
- **Snapshot copies graph**
This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

Autogrow	Displays whether the FlexVol volume will automatically grow in size when it is out of space.
Space Guarantee	<p>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:</p> <p>None No space guarantee is configured for the volume.</p> <p>File Full size of sparsely written files, for example LUNs, is guaranteed.</p> <p>Volume Full size of the volume is guaranteed.</p> <p>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.</p> <p>Note: The space guarantee is Partial when the volume is of type Data-Cache.</p>
Total Capacity	Displays the total capacity in the volume.
Data Capacity	<p>Displays the amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume.</p> <p>When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.</p>
Snapshot Reserve	<p>Displays the amount of space used by the Snapshot copies (used capacity) and amount of space available for Snapshot copies (free capacity) in the volume.</p> <p>When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the Snapshot copies (used capacity) and the amount of space that is available in the volume but cannot be used for making Snapshot copies (unusable capacity) because of aggregate capacity issues is displayed.</p> <p>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.</p>

**Volume
Thresholds**

Displays the following volume capacity thresholds:

- **Nearly Full Threshold**
Specifies the percentage at which a volume is nearly full.
- **Full Threshold**
Specifies the percentage at which a volume is full.

**Other
Details**

- **Autogrow Max Size**
Displays the maximum size up to which the FlexVol volume can automatically grow. The default value is 120% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.
- **Autogrow Increment Size**
Displays the increment size using which the size of the FlexVol volume increases every time the volume is automatically grown. The default is 5% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.
- **Qtree Quota Committed Capacity**
Displays the space reserved in the quotas.
- **Qtree Quota Overcommitted Capacity**
Displays the amount of space that can be used before the system generates the Volume Qtree Quota Overcommitted event.
- **Fractional Reserve**
Controls the size of the overwrite reserve. By default, the fractional reserve is set to 100, indicating that 100 percent of the required reserved space is reserved so that the objects are fully protected for overwrites. If the fractional reserve is less than 100 percent, the reserved space for all the space-reserved files in that volume is reduced to the fractional reserve percentage.
- **Snapshot Daily Growth Rate**
Displays the change (in percentage, and KB, MB, GB, and so on) that occurs every 24 hours in the Snapshot copies in the selected volume.
- **Snapshot Days to Full**
Displays the estimated number of days remaining before the space reserved for the Snapshot copies in the volume reaches the specified threshold.
The Snapshot Days to Full field displays a Not Applicable value when the growth rate of the Snapshot copies in the volume is zero, negative, or when there is insufficient data to calculate the growth rate.
- **Snapshot Autodelete**
Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.
- **Snapshot Copies**
Displays information about the Snapshot copies in the volume.
For volumes in a cluster running Data ONTAP 8.2 or later, the number of Snapshot copies in the volume is displayed as a link. Clicking the link opens the

Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

For volumes in a cluster running Data ONTAP 8.1.x, the View link is displayed. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

Volume Move

Displays the status of either the current or the last volume move operation that was performed on the volume and other details, such as the current phase of the volume move operation which is in progress, source aggregate, destination aggregate, start time, end time, and estimated end time.

It also displays the number of volume move operations that are performed on the selected volume. You can view more information about the volume move operations by clicking the Volume Move History link.

Efficiency tab

The Efficiency tab displays information about the space saved in the volumes by using storage efficiency features such as deduplication, compression, and FlexClone volumes:

- Deduplication**
 - **Enabled**
Specifies whether deduplication is enabled or disabled on a volume.
 - **Space Savings**
Displays the amount of space saved (in percentage, or KB, MB, GB, and so on) in a volume by using deduplication.
 - **Last Run**
Displays the time that has elapsed since the deduplication operation was last performed. Also specifies whether the deduplication operation was successful. If the time elapsed exceeds a week, the timestamp when the operation was performed is displayed.
 - **Mode**
Specifies whether the deduplication operation enabled on a volume is a manual, scheduled, or policy-based operation. If the mode is set to Scheduled, the operation schedule is displayed and if the mode is set to a policy, the policy name is displayed.
 - **Status**
Displays the current status of the deduplication operation. The status can be Idle, Initializing, Active, Undoing, Pending, Downgrading, or Disabled.
 - **Type**
Specifies the type of deduplication operation running on the volume. If the volume is in a SnapVault relationship, the type displayed is SnapVault. For any other volume, the type is displayed as Regular.

- Compression**
- **Enabled**
Specifies whether compression is enabled or disabled on a volume.
 - **Space Savings**
Displays the amount of space saved (in percentage, or KB, MB, GB, and so on) in a volume by using compression.

Configuration tab

The Configuration tab displays details about the selected volume, such as the export policy, RAID type, capacity and storage efficiency related features of the volume:

- Overview**
- **Full Name**
Displays the full name of the volume.
 - **Aggregate**
Displays the name of the aggregate that contains the volume.
 - **Storage Virtual Machine**
Displays the name of the Storage Virtual Machine (SVM) that contains the volume.
 - **Junction Path**
Displays the status of the path, which can be active or inactive. The path in the SVM to which the volume is mounted is also displayed. You can click the **History** link to view the last five changes to the junction path.
 - **Export policy**
Displays the name of the export policy that is created for the volume. You can click the link to view details about the export policies, authentication protocols, and access enabled on the volumes that belong to the SVM.
 - **Type**
Displays the type of the selected volume. The volume type can be Read-write, Load-sharing, Data-Protection, Data-cache, or Temporary.
 - **Style**
Displays the volume style, which is FlexVol.
 - **RAID Type**
Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, or RAID-DP.

- Capacity**
- **Thin Provisioning**
Displays whether thin provisioning is configured for the volume.
 - **Autogrow**
Displays whether the flexible volume grows in size automatically within an aggregate.
 - **Snapshot Autodelete**

Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.

- Quotas

Specifies whether the quotas are enabled for the volume.

Efficiency

- Deduplication

Specifies whether deduplication is enabled or disabled for the selected volume.

- Compression

Specifies whether compression is enabled or disabled for the selected volume.

Protection

- Snapshot Copies

Specifies whether automatic Snapshot copies are enabled or disabled.

Protection tab

The Protection tab displays protection details about the selected volume, such as lag information, relationship type, and topology of the relationship.

Summary Displays SnapMirror and SnapVault relationships properties for a selected volume. For any other relationship type, only the Relationship Type property is displayed. If a primary volume is selected, only the Managed and Local Snapshot copy Policy are displayed. Properties displayed for SnapMirror and SnapVault relationships include the following:

- Source Volume

Displays the name of the selected volume's source if the selected volume is a destination.

- Lag Status

Displays the update or transfer lag status for a protection relationship. The status can be Error, Warning, or Critical.

- Lag Duration

Displays the time by which the data on the mirror lags behind the source.

- Last Successful Update

Displays the date and time of the last successful protection update.

- Storage Service Member

Displays either Yes or No to indicate whether or not the volume belongs to and is managed by a storage service.

- Relationship Capability

Indicates the Data ONTAP capabilities available to the protection relationship. The relationship capability is either pre-8.2 or 8.2 and later. A relationship capability of pre-8.2 means that relationships have not been upgraded to Data ONTAP 8.2 on both the destination and the source clusters, and cannot take advantage of new or improved protection features available in Data ONTAP 8.2 and later. A relationship

capability of 8.2 and later means the destination and source clusters are using Data ONTAP 8.2 or later, and can take advantage of improved protection features.

- **Protection Service**
Displays the name of the protection service if the relationship is managed by a protection partner application.
- **Relationship Type**
Displays any relationship type, including SnapMirror or SnapVault.
- **Relationship State**
Displays the state of the SnapMirror or SnapVault relationship. The state can be Uninitialized, SnapMirrored, or Broken-Off. If a source volume is selected, the relationship state is not applicable and is not displayed.
- **Transfer Status**
Displays the transfer status for the protection relationship. The transfer status can be one of the following:
 - **Idle**
Transfers are enabled and no transfer is in progress.
 - **Transferring**
SnapMirror transfers are enabled and a transfer is in progress.
 - **Checking**
The destination volume is undergoing a diagnostic check and no transfer is in progress. This applies only to SnapMirror relationships that have the relationship-control-plane field set to v1.
 - **Quiescing**
A SnapMirror transfer is in progress. Additional transfers are disabled.
 - **Quiesced**
SnapMirror transfers are disabled. No transfer is in progress.
 - **Queued**
SnapMirror transfers are enabled. No transfers are in progress.
 - **Preparing**
SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.
 - **Finalizing**
SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.
 - **Aborting**
SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.
- **Max Transfer Rate**
Displays the maximum transfer rate for the relationship. The maximum transfer rate can be a numerical value in either kilobytes per second (Kbps), Megabytes

per second (Mbps), Gigabytes per second (Gbps), or Terabytes per second (TBps). If No Limit is displayed, the baseline transfer between relationships is unlimited.

- **SnapMirror Policy**

Displays the protection policy for the volume. DPDefault indicates the default SnapMirror protection policy, and XDPDefault indicates the default SnapVault policy. You can click the policy name to view details associated with that policy, including the following information:

- Transfer priority
- Ignore access time setting
- Tries limit
- Comments
- SnapMirror labels
- Retention settings
- Actual Snapshot copies
- Preserve Snapshot copies
- Retention warning threshold
- Snapshot copies with no retention settings

In a cascading SnapVault relationship where the source is a data protection (DP) volume, only the rule “sm_created” applies.

- **Update Schedule**

Displays the SnapMirror schedule assigned to the relationship. Positioning your cursor over the information icon displays the schedule details.

- **Local Snapshot Policy**

Displays the Snapshot copy policy for the volume. The policy is Default, None, or any name given to a custom policy.

Views

Displays the protection topology of the selected volume. The topology includes graphical representations of all volumes that are related to the selected volume. The selected volume is indicated by a dark gray border, and lines between volumes in the topology indicate the protection relationship type. Double lines specify a SnapMirror relationship, and a single line specifies a SnapVault relationship. The direction of the relationships in the topology are displayed from left to right, with the source of each relationship on the left and the destination on the right.

Right-clicking a volume displays a menu from which you can choose either to protect the volume or restore data to it.

Right-clicking a relationship displays a menu from which you can choose to either edit, abort, quiesce, break, remove, or resume a relationship. The menus will not display in the following instances:

- If RBAC settings do not allow this action, for example, if you have only operator privileges

- When the volume ID is unknown, for example, when you have a intercluster relationship and the destination cluster has not yet been discovered
- When the volume is a Data ONTAP 8.1 cluster volume

Clicking another volume in the topology selects and displays information for that

volume. A question mark () in the upper-left corner of a volume indicates that either the volume is missing or that it has not yet been discovered. It might also indicate that the capacity information is missing. Positioning your cursor over the question mark displays additional information, including suggestions for remedial action.

The topology displays information about volume capacity, lag, Snapshot copies, and last successful data transfer if it conforms to one of several common topology templates. If a topology does not conform to one of those templates, information about volume lag and last successful data transfer is displayed in a relationship table under the topology. In that case, the highlighted row in the table indicates the selected volume, and, in the topology view, bold lines with a blue dot indicate the relationship between the selected volume and its source volume.

Topology views include the following information:

- **Capacity**
Displays the total amount of capacity used by the volume. Positioning your cursor over a volume in the topology displays the current warning and critical threshold settings for that volume in the Current Threshold Settings dialog box. You can also edit the threshold settings by clicking the **Edit Thresholds** link in the Current Threshold Settings dialog box. Clearing the **Capacity** check box hides all capacity information for all volumes in the topology.
- **Lag**
Displays the lag duration and the lag status of the incoming protection relationships. Clearing the **Lag** check box hides all lag information for all volumes in the topology. When the **Lag** check box is dimmed, then the lag information for the selected volume is displayed in the relationship table below the topology, as well as the lag information for all related volumes.
- **Snapshot**
Displays the number of Snapshot copies available for a volume. Clearing the **Snapshot** check box hides all Snapshot copy information for all volumes in the topology. Clicking a Snapshot copy icon () displays the Snapshot copy list for a volume. The Snapshot copy count displayed next to the icon is updated every 15 minutes; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon. If you are running Data ONTAP 8.1, the Snapshot copy count is not displayed in the topology.
- **Last Successful Transfer**

Displays the amount, duration, time, and date of the last successful data transfer. When the **Last Successful Transfer** check box is dimmed, then the last successful transfer information for the selected volume is displayed in the relationship table below the topology, as well as the last successful transfer information for all related volumes.

History Displays in a graph the history of incoming SnapMirror and SnapVault protection relationships for the selected volume when you are using Data ONTAP 8.2 or later relationship capabilities. No historical data is collected for relationship capabilities earlier than Data ONTAP 8.2. There are three history graphs available: incoming relationship transfer duration, incoming relationship lag size, and incoming relationship transferred size. History information is displayed only when you select a destination volume. If you select a primary volume, the graphs are empty, and the message `No data found` is displayed.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if large amounts of data are being transferred at the same time of the day or week, or if the lag warning or lag error threshold is consistently being breached, you can take the appropriate action.

History graphs display the following information:

- | | |
|---------------------------------------|---|
| Relationship Transfer Duration | Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum transfer duration reached in the duration period shown in the x axis. You can view the details of specific points on the graph by positioning your cursor over the area of interest. |
| Relationship Lag Duration | Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum lag duration reached in the duration period shown in the x axis. The horizontal orange line on the graph depicts the lag error threshold, and the horizontal yellow line depicts the lag warning threshold. Positioning your cursor over these lines displays the threshold setting. The horizontal blue line depicts the lag duration. You can view the details for specific points on the graph by positioning your cursor over an area of interest. |
| Relationship Transferred Size | Displays bytes, kilobytes, megabytes, and so on, on the vertical (y) axis depending on the transfer size, and displays days, months, or years on the horizontal (x) axis depending on the selected time period. The upper value on the y axis indicates the maximum transfer size reached in the duration period shown in the x axis. You |

can view the details for specific points on the graph by positioning your cursor over an area of interest.

History area

The History area displays graphs that provide information about the capacity and space reservations of the selected volume.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends—for example, if the volume usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Volume Capacity Used	Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.
Volume Capacity Used vs Total	Displays the trend in how volume capacity is used based on the usage history, as well as the used capacity, total capacity, and details of the space savings from deduplication and compression, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.
Volume Capacity Used (%)	Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.
Snapshot Capacity Used (%)	Displays the Snapshot reserve and Snapshot warning threshold as line graphs, and the capacity used by the Snapshot copies as an area graph, in percentage, on the vertical (y) axis. The Snapshot overflow is represented with different colors. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a

month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Snapshot Reserve legend, the Snapshot Reserve graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity	Displays the severity of the event.
Event	Displays the event name.
Triggered Time	Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp when the event was generated is displayed.

Related Devices pane

The Related Devices pane enables you to view and navigate to the SVMs, aggregates, qtrees, LUNs, and Snapshot copies that are related to the volume:

Storage Virtual Machine	Displays the capacity and the health status of the SVM that contains the selected volume.
Aggregate	Displays the capacity and the health status of the aggregate that contains the selected volume.
Volumes in the Aggregate	Displays the number and capacity of all the volumes that belong to the parent aggregate of the selected volume. The health status of the volumes is also displayed, based on the highest severity level. For example, if an aggregate contains ten volumes, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.
Qtrees	Displays the number of qtrees that the selected volume contains and the capacity of qtrees with quota that the selected volume contains. The capacity of the qtrees with quota is displayed in relation to the volume data capacity. The health status of the qtrees is also displayed, based on the highest severity level. For example, if a volume has ten qtrees, five with Warning status and the remaining five with Critical status, then the status displayed is Critical.
NFS Exports	Displays the number and status of the NFS exports associated with the volume.
CIFS Shares	Displays the number and status of the CIFS shares.
LUNs	Displays the number and total size of all the LUNs in the selected volume. The health status of the LUNs is also displayed, based on the highest severity level.
User and Group Quotas	Displays the number and status of the user and user group quotas associated with the volume and its qtrees.

- FlexClone Volumes** Displays the number and capacity of all the cloned volumes of the selected volume. The number and capacity are displayed only if the selected volume contains any cloned volumes.
- Parent Volume** Displays the name and capacity of the parent volume of a selected FlexClone volume. The parent volume is displayed only if the selected volume is a FlexClone volume.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected volume. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Annotations pane

The Annotations pane enables you to view the annotation configured for the selected volume.

Related tasks

- [Performing diagnostic actions for volume offline conditions](#) on page 48
- [Performing suggested remedial actions for a full volume](#) on page 69

Storage Virtual Machine details page

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

Note: You can monitor only data SVMs.

- [Command buttons](#) on page 117
- [Health tab](#) on page 117
- [Capacity tab](#) on page 118
- [Configuration tab](#) on page 120
- [LIFs tab](#) on page 122
- [Qtrees tab](#) on page 122
- [User and Group Quotas tab](#) on page 124
- [NFS Exports tab](#) on page 126
- [CIFS Shares tab](#) on page 127
- [SAN tab](#) on page 128
- [Data Policy tab](#) on page 129
- [Related Devices pane](#) on page 130

- [Related Alerts pane](#) on page 131
- [Annotations pane](#) on page 131

Command buttons

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

Actions

- **Add Alert**
Enables you to add an alert to the selected SVM.
- **Edit Thresholds**
Enables you to edit the SVM thresholds.

Note: This button is enabled only for SVM with Infinite Volume.

View Storage Virtual Machines

Enables you to navigate to the Storage Virtual Machines page.

Health tab

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

You can click the graph of an object to view the filtered list of objects. For example, you can click the volume capacity graph that displays warnings to view the list of volumes that have capacity issues with severity as warning.

Availability Displays, as a graph, the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the SVM. For example, information is displayed about the NAS LIFs and the SAN LIFs that are down and volumes that are offline.

You can also view information about the related protocols and services that are currently running, and the number and status of NFS exports and CIFS shares.

If the selected SVM is an SVM with Infinite Volume, you can view availability details about the Infinite Volume.

Capacity Displays, as a graph, the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted

the capacity of data in the SVM. For example, information is displayed about aggregates that are likely to breach the set threshold values.

If the selected SVM is an SVM with Infinite Volume, you can view capacity details about the Infinite Volume.

Capacity tab

The Capacity tab displays detailed information about the data capacity of the selected SVM.

The following information is displayed for an SVM with FlexVol volume:

- Capacity** The Capacity area displays details about the used and available capacity allocated from all volumes:
- **Total Capacity**
Displays the total capacity (in MB, GB, and so on) of the SVM.
 - **Used**
Displays the space used by data in the volumes that belong to the SVM.
 - **Guaranteed Available**
Displays the guaranteed available space for data that is available for volumes in the SVM.
 - **Unguaranteed**
Displays the available space remaining for data that is allocated for thinly provisioned volumes in the SVM.

Volumes with Capacity Issues The Volumes with Capacity Issues list displays, in tabular format, details about the volumes that have capacity issues:

- **Status**
Indicates that the volume has a capacity-related issue of a certain severity. You can move the pointer over the status to view more information about the capacity-related event or events generated for the volume.
If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use the **View Details** button to view more information about the event.
If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

Note: A volume can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a volume has two events with severities of Error and Warning, only the Error severity is displayed.

- **Volume**
Displays the name of the volume.
- **Used Data Capacity**
Displays, as a graph, information about the volume capacity usage (in percentage).
- **Days to Full**
Displays the estimated number of days remaining before the volume reaches full capacity.
- **Thin Provisioned**
Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.
- **Aggregate**
Displays the name of the aggregate that contains the volume.

The following information is displayed for an SVM with Infinite volume:

- | | |
|----------------------|---|
| Capacity | <p>Displays the following capacity-related details:</p> <ul style="list-style-type: none"> • Percentage of used and free data capacity • Percentage of used and free Snapshot capacity • Snapshot Overflow
Displays the data space that is consumed by the Snapshot copies. • Used
Displays the space used by data in the SVM with Infinite Volume. • Warning
Indicates that the space in the SVM with Infinite Volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated. • Error
Indicates that the space in the SVM with Infinite Volume is full. If this threshold is breached, the Space Full event is generated. |
| Other Details | <ul style="list-style-type: none"> • Total Capacity
Displays the total capacity in the SVM with Infinite Volume. • Data Capacity
Displays used data capacity, available data capacity, and Snapshot overflow capacity details of the SVM with Infinite Volume. • Snapshot Reserve
Displays the used and free details of the Snapshot reserve. |

- **System Capacity**
Displays the used system capacity and available system capacity in the SVM with Infinite Volume.
- **Thresholds**
Displays the nearly full and full thresholds of the SVM with Infinite Volume.

Storage Class Capacity Details Displays information about the capacity usage in your storage classes. This information is displayed only if you have configured storage classes for your SVM with Infinite Volume.

Storage Virtual Machine Storage Class Thresholds Displays the following thresholds (in percentage) of your storage classes:

- **Nearly Full Threshold**
Specifies the percentage at which a storage class in an SVM with Infinite Volume is considered to be nearly full.
- **Full Threshold**
Specifies the percentage at which the storage class in an SVM with Infinite Volume is considered full.
- **Snapshot Usage Limit**
Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.

Configuration tab

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

Overview

- **Cluster**
Displays the name of the cluster to which the SVM belongs.
- **Allowed Volume Type**
Displays the type of volumes that can be created in the SVM. The type can be InfiniteVol or FlexVol.
- **Root Volume**
Displays the name of the root volume of the SVM.
- **Allowed Protocols**
Displays the type of protocols that can be configured on the SVM. Also, indicates if a protocol is up (), down (), or is not configured ().

Data LIFs

- **NAS**
Displays the number of NAS LIFs that are associated with the SVM. Also, indicates if the LIFs are up () or down ().

- **SAN**
Displays the number of SAN LIFs that are associated with the SVM. Also, indicates if the LIFs are up () or down ().
- **Junction Path**
Displays the path on which the Infinite Volume is mounted. Junction path is displayed for an SVM with Infinite Volume only.
- **Storage Classes**
Displays the storage classes associated with the selected SVM with Infinite Volume. Storage classes are displayed for an SVM with Infinite Volume only.

Management LIFs

- **Availability**
Displays the number of management LIFs that are associated with the SVM. Also, indicates if the management LIFs are up () or down ().

Policies

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

Services

- **Type**
Displays the type of service that is configured on the SVM. The type can be Domain Name System (DNS) or Network Information Service (NIS).
- **State**
Displays the state of the service, which can be Up (), Down (), or Not Configured ().
- **Domain Name**
Displays the fully qualified domain names (FQDNs) of the DNS server for the DNS services or NIS server for the NIS services. When the NIS server is enabled, the active FQDN of the NIS server is displayed. When the NIS server is disabled, the list of all the FQDNs are displayed.
- **IP Address**
Displays the IP addresses of the DNS or NIS server. When the NIS server is enabled, the active IP address of the NIS server is displayed. When the NIS server is disabled, the list of all the IP addresses are displayed.

LIFs tab

The LIFs tab displays details about the data LIFs that are created on the selected SVM:

LIF	Displays the name of the LIF that is created on the selected SVM.
Operational Status	Displays the operational status of the LIF, which can be Up () , Down () , or Unknown (). The operational status of a LIF is determined by the status of its physical ports.
Administrative Status	Displays the administrative status of the LIF, which can be Up () , Down () , or Unknown (). The administrative status of a LIF is controlled by the storage administrator to make changes to the configuration or for maintenance purposes. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.
IP Address / WWPN	Displays the IP address for Ethernet LIFs and the World Wide Port Name (WWPN) for FC LIFs.
Protocols	Displays the list of data protocols that are specified for the LIF, such as CIFS, NFS, iSCSI, FC/FCoE, and FlexCache. For Infinite Volume, the SAN protocols are not applicable.
Role	Displays the LIF role. The roles can be Data or Management.
Home Port	Displays the physical port to which the LIF was originally associated.
Port Set	Displays the port set to which the LIF is mapped.
Current Port	Displays the physical port to which the LIF is currently associated. If the LIF is migrated, the current port might be different from the home port.
Failover Policy	Displays the failover policy that is configured for the LIF. For NFS, CIFS, and FlexCache LIFs, the default failover policy is Next Available. Failover policy is not applicable for FC and iSCSI LIFs.
Routing Groups	Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.
Failover Group	Displays the name of the failover group.

Qtrees tab

The Qtrees tab displays details about qtrees and their quotas:

Note: The Qtrees tab is not displayed for an SVM with Infinite Volume.

Status

Displays the current status of the qtree. The status can be Critical () , Error () , Warning () , or Normal () .

You can move the pointer over the status icon to view more information about the event or events generated for the qtree.

If the status of the qtree is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the qtree is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

Note: A qtree can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a qtree has two events with severities of Error and Warning, only the Error severity is displayed.

Qtree

Displays the name of the qtree.

Volume

Displays the name of the volume that contains the qtree.

You can move the pointer over the volume name to view more information about the volume.

Quota Set

Indicates whether a quota is enabled or disabled on the qtree.

Disk Used %

Displays the percentage of disk space used. If a disk hard limit is set, this value is based on the disk hard limit. If the quota is set without a disk hard limit, the value is based on the volume data space. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then “Not applicable” is displayed.

Disk Hard Limit

Displays the maximum amount of disk space allocated for the qtree. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

Disk Soft Limit

Displays the amount of disk space allocated for the qtree before a warning event is generated. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

Disk Threshold	Displays the threshold value set on the disk space. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk threshold limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.
Files Used %	Displays the percentage of files used in the qtree. If the file hard limit is set, this value is based on the file hard limit. No value is displayed if the quota is set without a file hard limit. The value is displayed as “Not applicable” if the quota is not set or if quotas are off on the volume to which qtree belongs.
File Hard Limit	Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.
File Soft Limit	Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

User and Group Quotas tab

Displays details about the user and user group quotas for the selected SVM. You can view information such as the status of the quota, name of the user or user group, soft and hard limits set on the disks and files, amount of disk space and number of files used, and the disk threshold value. You can also change the email address associated with a user or user group.

Edit Email Address Opens the Edit Email Address dialog box, which displays the current email address of the selected user or user group. You can modify the email address. If the **Edit Email Address** field is blank, the default rule is used to generate an email address for the selected user or user group.

If more than one user has the same quota, the names of the users are displayed as comma-separated values. Also, the default rule is not used to generate the email address; therefore, you must provide the required email address for notifications to be sent.

Configure Email Rules Enables you to create or modify rules to generate an email address for the user or user group quotas that are configured on the SVM. A notification is sent to the specified email address when there is a quota breach.

Status Displays the current status of the quota. The status can be Critical () , Warning () , or Normal () .

You can move the pointer over the status icon to view more information about the event or events generated for the quota.

If the status of the quota is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the quota is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

Note: A quota can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a quota has two events with severities of Error and Warning, only the Error severity is displayed.

User or Group	<p>Displays the name of the user or user group. If more than one user has the same quota, the names of the users are displayed as comma-separated values.</p> <p>The value is displayed as “Unknown” when Data ONTAP does not provide a valid user name because of SecD errors.</p>
Type	Specifies if the quota is for a user or a user group.
Volume or Qtree	<p>Displays the name of the volume or qtree on which the user or user group quota is specified.</p> <p>You can move the pointer over the name of the volume or qtree to view more information about the volume or qtree.</p>
Disk Used %	Displays the percentage of disk space used. The value is displayed as “Not applicable” if the quota is set without a disk hard limit.
Disk Soft Limit	Displays the amount of disk space allocated for the quota before a warning event is generated. The value is displayed as “Unlimited” if the quota is set without a disk soft limit. By default, this column is hidden.
Disk Hard Limit	Displays the maximum amount of disk space allocated for the quota. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” if the quota is set without a disk hard limit.
Disk Threshold	Displays the threshold value set on the disk space. The value is displayed as “Unlimited” if the quota is set without a disk threshold limit. By default, this column is hidden.

Files Used %	Displays the percentage of files used in the qtree. The value is displayed as “Not applicable” if the quota is set without a file hard limit.
File Hard Limit	Displays the hard limit for the number of files permitted on the quota. The value is displayed as “Unlimited” if the quota is set without a file hard limit.
File Soft Limit	Displays the soft limit for the number of files permitted on the quota. The value is displayed as “Unlimited” if the quota is set without a file soft limit. By default, this column is hidden.
Email Address	Displays the email address of the user or user group to which notifications are sent when there is a breach in the quotas.

NFS Exports tab

The NFS Exports tab displays information about the NFS exports such as its status, the path associated with the volume (Infinite Volumes or FlexVol volumes), access levels of clients to the NFS exports, and the export policy defined for the volumes that are exported. NFS Export will not be displayed for the following conditions: if the volume is not mounted or if the protocols associated with the export policy for the volume do not contain NFS.

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

- **UNIX (NFS clients)**
Files and directories in the volume have UNIX permissions.
- **Unified**
Files and directories in the volume have a unified security style.
- **NTFS (CIFS clients)**
Files and directories in the volume have Windows NTFS permissions.
- **Mixed**
Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

UNIX Permission	Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.
Export Policy	Displays the rules that define the access permission for volumes that are exported. You can click the link to view details about the rules associated with the export policy such as the authentication protocols and the access permission.

CIFS Shares tab

Displays information about the CIFS shares on the selected SVM. You can view information such as the status of the CIFS share, share name, path associated with the SVM, the status of the junction path of the share, containing object, state of the containing volume, security data of the share, and export policies defined for the share. You can also determine whether an equivalent NFS path for the CIFS share exists.

View User Mapping	Launches the User Mapping dialog box. You can view the details of user mapping for the SVM.
Show ACL	Launches the Access Control dialog box for the share. You can view user and permission details for the selected share.

Status	Displays the current status of the share. The status can be Normal () or Error ()
Share Name	Displays the name of the CIFS share.
Path	Displays the junction path on which the share is created.
Junction Path Active	Displays whether the path to access the share is active or inactive.
Containing Object	Displays the name of the containing object to which the share belongs. The containing object can be a volume or a qtree.

By clicking the link, you can view details about the containing object in the respective Details page. If the containing object is a qtree, links are displayed for both qtree and volume.

Volume State	<p>Displays the state of the volume that is being exported. The state can be Offline, Online, or Restricted.</p> <ul style="list-style-type: none"> • Offline Read or write access to the volume is not allowed. • Online Read and write access to the volume is allowed. • Restricted Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
Security	<p>Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.</p> <ul style="list-style-type: none"> • UNIX (NFS clients) Files and directories in the volume have UNIX permissions. • NTFS (CIFS clients) Files and directories in the volume have Windows NTFS permissions. • Mixed Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.
Export Policy	<p>Displays the name of the export policy applicable on the share. If an export policy is not specified for the , the value is displayed as Not Enabled.</p> <p>You can click the link to view details about the rules associated with the export policy, such as access protocols and permissions. The link is disabled if the export policy is disabled for the selected SVM.</p>
NFS Equivalent	<p>Specifies whether there is an NFS equivalent for the share.</p>

SAN tab

Displays details about LUNs, initiator groups, and initiators for the selected SVM. By default, the LUNs view is displayed. You can view details about the initiator groups in the Initiator Groups tab and details about initiators in the Initiators tab.

LUNs	<p>Displays details about the LUNs that belong to the selected SVM. You can view information such as the LUN name, LUN state (online or offline), the name of the file system (volume or qtree) that contains the LUN, the type of host operating system, the total data capacity and serial number of the LUN. You can also view information</p>
-------------	---

whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group.

You can also view the initiator groups and initiators that are mapped to the selected LUN.

Initiator Groups Displays details about initiator groups. You can view details such as the name of the initiator group, the access paths to which the initiator group is connected (one, many, or no paths), the type of host operating system used by all of the initiators in the group, and the supported protocol. You can view if initiator groups are mapped to all the LIFs or specific LIFs through a port set. When you click the count link in the Mapped LIFs column, the LIFs that are mapped through the target portal are not displayed. The total number of initiators and LUNs that are mapped to an initiator group is displayed.

You can also view the LUNs and initiators that are mapped to the selected initiator group.

Initiators Displays the name and type of the initiator and the total number of initiator groups mapped to this initiator for the selected SVM.

You can also view the LUNs and initiator groups that are mapped to the selected initiator group.

Data Policy tab

The Data Policy tab enables you to create, modify, activate, or delete one or more rules in a data policy. You can also import the data policy into the Unified Manager database and export the data policy to your computer:

Note: The Data Policy tab is displayed only for SVMs with Infinite Volume.

Rules list Displays the list of rules. By expanding the rule, you can view the corresponding matching criteria of the rule and the storage class where the content is placed based on the rule.

The default rule is the last rule in the list. You cannot change the order of the default rule.

- **Matching Criteria**

Displays the conditions for the rule. For example, a rule can be “File path starts with /eng/nightly”.

Note: The file path must always start with a junction path.

- **Content Placement**

Displays the corresponding storage class for the rule.

Rule Filter Enables you to filter rules associated with a specific storage class listed in the list.

Action buttons

- **Create**
Opens the Create Rule dialog box, which enables you to create a new rule for your data policy.
- **Edit**
Opens the Edit Rule dialog box, which enables you to modify rule properties such as directory paths, file types, and owners.
- **Delete**
Deletes the selected rule.
- **Move Up**
Moves the selected rule up in the list. However, you cannot move the default rule up in the list.
- **Move Down**
Moves the selected rule down the list. However, you cannot move the default rule down the list.
- **Activate**
Activates the rules and changes made to the data policy in the SVM with Infinite Volume.
- **Reset**
Resets all changes made to the data policy configuration.
- **Import**
Imports a data policy configuration from a file.
- **Export**
Exports a data policy configuration to a file.

Related Devices area

The Related Devices area enables you to view and navigate to the LUNs, CIFS shares, and the user and user group quotas that are related to the qtree:

- LUNs** Displays the total number of the LUNs associated with the selected qtree.
- CIFS Shares** Displays the total number of CIFS shares associated with the selected qtree.
- User and Group Quotas** Displays the total number of the user and user group quotas associated with the selected qtree. The health status of the user and user group quotas is also displayed, based on the highest severity level.

Related Devices pane

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

- Cluster** Displays the health status of the cluster to which the SVM belongs.

Aggregates	Displays the number of aggregates that belong to the selected SVM. The health status of the aggregates is also displayed, based on the highest severity level. For example, if an SVM contains ten aggregates, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.
Assigned Aggregates	Displays the number of aggregates that are assigned to an SVM. The health status of the aggregates is also displayed, based on the highest severity level.
Volumes	Displays the number and capacity of the volumes that belong to the selected SVM. The health status of the volumes is also displayed, based on the highest severity level.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected SVM. You can also add an alert by clicking the **Add Alert** link or edit an existing alert by clicking the alert name.

Annotations pane

The Annotations pane enables you to view the annotation for the selected SVM.

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 131
- [Health tab](#) on page 132
- [Capacity tab](#) on page 133
- [Configuration tab](#) on page 134
- [LIFs tab](#) on page 135
- [Nodes tab](#) on page 136
- [Disks tab](#) on page 138
- [Related Devices pane](#) on page 139
- [Related Alerts pane](#) on page 139
- [Annotations pane](#) on page 140

Command buttons

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

Actions

- **Add Alert**
Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.
- **Edit Cluster**
Opens the Edit Cluster dialog box, which enables you to edit the settings of the selected cluster.
- **Rediscover**
Initiates a manual refresh of a cluster, which enables Unified Manager to discover recent changes to the cluster. If Unified Manager is paired with OnCommand Workflow Automation, this operation also reacquires WFA's cached data.
After the rediscovery operation is initiated, a link to the associated job details is displayed to enable tracking the job status.

View Clusters Enables you to navigate to the Clusters page.

Health tab

The Health tab displays detailed information about data availability and data capacity issues of various cluster objects such as nodes, Storage Virtual Machines (SVMs), and aggregates. Availability issues are related to the data serving capability of the cluster objects. Capacity issues are related to the data storing capability of the cluster objects.

You can click the graph of an object to view the filtered list of the objects. For example, you can click the SVM capacity graph that displays warnings to view the filtered list of SVMs. The list contains SVMs that have volumes or qtrees that have capacity issues with severity as warning. You can also click the SVMs availability graph that displays warnings to view the list of SVMs that have availability issues with severity as warning.

Availability Displays, as a graph, the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the cluster. For example, information is displayed about disk shelves that are down and aggregates that are offline.

Note: The data displayed for the SFO bar graph is based on the HA state of the nodes. The data displayed for all other bar graphs is calculated based on the events generated.

Capacity Displays, as a graph, the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted

the capacity of data in the cluster. For example, information is displayed about aggregates that are likely to breach the set threshold values.

Capacity tab

The Capacity tab displays detailed information about the capacity of the selected cluster:

Capacity	<p>The Capacity area displays details about the used and available capacity from all allocated aggregates:</p> <ul style="list-style-type: none"> • Total Capacity Displays the total capacity (in MB, GB, and so on) of the cluster. This does not include the capacity that is assigned for parity. • Used Displays the capacity (in MB, GB, and so on) that is used by data. This does not include the capacity that is used for parity, right-sizing, and reservation. • Available Displays the capacity (in MB, GB, and so on) available for data. • Provisioned Displays the capacity (in MB, GB, and so on) that is provisioned for all the underlying volumes. • Spares Displays the storable capacity (in MB, GB, and so on) available for storage in all the spare disks. • SSD Tier Displays the total space of the solid-state disks (SSDs) that are added to the cluster.
Capacity Breakout by Disk Type	<p>The Capacity Breakout by Disk Type area displays detailed information about the disk capacity of various types of disks in the cluster. By clicking the disk type, you can view more information about the disk type from the Disks tab.</p> <ul style="list-style-type: none"> • Total Usable Capacity Displays the distribution of available, used, and spare capacity disks. The dotted line represents the total size of the spare disk capacity available for a disk type. • Unassigned Disks Displays the number of unassigned disks in the cluster. • Cache Displays the total size of spare disk capacity.
Aggregates with Capacity Issues list	<p>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:</p> <ul style="list-style-type: none"> • Status

Indicates that the aggregate has a capacity-related issue of a certain severity. You can move the pointer over the status to view more information about the event or events generated for the aggregate.

If the status of the aggregate is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event.

If the status of the aggregate is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

Note: An aggregate can have multiple capacity-related events of the same severity or different severities. However, only the highest severity is displayed. For example, if an aggregate has two events with severities of Error and Critical, only the Critical severity is displayed.

- **Aggregate**
Displays the name of the aggregate.
- **Used Data Capacity**
Displays, as a graph, information about the aggregate capacity usage (in percentage).
- **Days to Full**
Displays the estimated number of days remaining before the aggregate reaches full capacity.

Configuration tab

The Configuration tab displays details about the selected cluster, such as the IP address, serial number, contact, and location information of the cluster:

- Overview**
- **Management LIF**
Displays the cluster-management LIF that Unified Manager uses to connect to the cluster. The operational status of the LIF is also displayed.
 - **Host Name or IP Address**
Displays the FQDN, short name, or the IP address of the cluster-management LIF that Unified Manager uses to connect to the cluster.
 - **FQDN**
Displays the Fully Qualified Domain Name (FQDN) of the cluster.
 - **OS Version**

Displays the Data ONTAP version that the cluster is running. If the nodes in the cluster are running different versions of Data ONTAP, then the earliest Data ONTAP version is displayed.

- **Serial Number**
Displays the serial number of the cluster.
- **Contact**
Displays the contact information of the cluster.
- **Location**
Displays the location of the cluster.

Nodes

- **Availability**
Displays the number of nodes that are up () or down () in the cluster.
- **OS Versions**
Displays the Data ONTAP versions that the nodes are running. Also, displays the number of nodes running a particular version of Data ONTAP. For example, 8.2 (2), 8.1 (1) specifies that two nodes are running Data ONTAP 8.2 and one node is running Data ONTAP 8.1.

Storage Virtual Machines

- **Availability**
Displays the number of SVMs that are up () or down () in the cluster.

LIFs

- **Availability**
Displays the number of non-data LIFs that are up () or down () in the cluster.
- **Cluster-Management LIFs**
Displays the number of cluster-management LIFs.
- **Node-Management LIFs**
Displays the number of node-management LIFs.
- **Cluster LIFs**
Displays the number of cluster LIFs.
- **Intercluster LIFs**
Displays the number of intercluster LIFs.

Protocols

- **Data Protocols**
Displays the list of licensed data protocols that are enabled for the cluster. The data protocols include iSCSI, CIFS, NFS, and FC and FCoE.

LIFs tab

The LIFs tab displays details about all the non-data LIFs that are created on the selected cluster:

LIF	Displays the name of the LIF that is created on the selected cluster.
Operational Status	Displays the operational status of the LIF, which can be Up () , Down () , or Unknown (). The operational status of a LIF is determined by the status of its physical ports.
Administrative Status	Displays the administrative status of the LIF, which can be Up () , Down () , or Unknown (). You can control the administrative status of a LIF when you make changes to the configuration or during maintenance. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.
IP Address	Displays the IP address of the LIF.
Role	Displays the role of the LIF. Possible roles are Cluster-Management LIFs, Node-Management LIFs, Cluster LIFs, and Intercluster LIFs.
Home Port	Displays the physical port to which the LIF was originally associated.
Current Port	Displays the physical port to which the LIF is currently associated. After LIF migration, the current port might be different from the home port.
Failover Policy	Displays the failover policy that is configured for the LIF.
Routing Groups	Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.
Failover Group	Displays the name of the failover group.

Nodes tab

The Nodes tab displays information about nodes in the selected cluster. You can view detailed information about the HA pairs, disk shelves, and ports:

HA Details	<p>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:</p> <ul style="list-style-type: none"> • Green: The node is in a working condition. • Yellow: The node has taken over the partner node or the node is facing some environmental issues. • Red: The node is down.
-------------------	---

You can view information about the availability of the HA pair and take required action to prevent any risks. For example, in the case of a possible takeover operation, the following message is displayed: Storage failover possible

You can view a list of the events related to the HA pair and its environments, such as fans, power supplies, NVRAM battery, flash cards, service processor, and connectivity of disk shelves. You can also view the time when the events were triggered.

You can view other node-related information, such as the model number and the serial number.

If there are single-node clusters, you can also view details about the nodes.

Disk Shelves Displays information about the disk shelves in the HA pair.

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

Shelf ID	Displays the ID of the shelf where the disk is located.
Component Status	Displays environmental details of the disk shelves, such as power supplies, fans, temperature sensors, current sensors, disk connectivity, and voltage sensors. The environmental details are displayed as icons in the following colors: <ul style="list-style-type: none"> • Green: The environmental components are in working properly. • Grey: No data is available for the environmental components. • Red: Some of the environmental components are down.
State	Displays the state of the disk shelf. The possible states are Offline, Online, No status, Initialization required, Missing, and Unknown.
Model	Displays the model number of the disk shelf.
Unique ID	Displays the unique identifier of the disk shelf.
Firmware Version	Displays the firmware version of the disk shelf.

Ports Displays information about the associated FC, FCoE, and Ethernet port details. You can view details about the ports and the associated LIFs by clicking the port icons.

You can also view the events generated for the ports.

You can view the following port details:

- **Port ID**
Displays the name of the port. For example, the port names can be e0M, e0a, and e0b.
- **Role**
Displays the role of the port. The possible roles are Cluster, Data, Intercluster, Node-Management, and Undefined.
- **Type**
Displays the physical layer protocol used for the port. The possible types are Ethernet, Fibre Channel, and FCoE.

- **WWPN**
Displays the World Wide Port Name (WWPN) of the port.
- **Firmware Rev**
The firmware revision of the FC/FCoE port.
- **Status**
Displays the current state of the port. The possible states are Online, Offline, or Unknown ().

You can view the port-related events from the Events list. You can also view the associated LIF details, such as LIF name, operational status, the IP address or the WWPN, protocols, name of the SVM associated with the LIF, current port, failover policy and failover group.

Disks tab

The Disks tab displays details about the disks in the selected cluster. You can view disk-related information such as the number of used disks, spare disks, broken disks, and unassigned disks. You can also view other details such as the disk name, disk type, and the owner node of the disk:

Disk Pool Summary	Displays the number of disks, which are categorized by effective types (FCAL, SAS, SATA, MSATA, SSD, and Array LUN), and the state of the disks. You can also view other details, such as the number of broken disks, spare disks, and unassigned disks. If you click the effective disk type count link, disks of the selected state and effective type are displayed. For example, if you click the count link for the disk state Broken and effective type SAS, all disks with the disk state Broken and effective type SAS are displayed.
Disk	Displays the name of the disk.
RAID Groups	Displays the name of the RAID group.
Owner Node	Displays the name of the node to which the disk belongs. If the disk is unassigned, no value is displayed in this column.
State	Displays the state of the disk—Broken, Aggregate, Spare, Unknown, or Unassigned. By default, this column is sorted to display the states in the following order: Broken, Spare, Aggregate, Unknown, and Unassigned.
Position	Displays the position of the disk based on its container type—for example, Copy, Data, or Parity. By default, this column is hidden.
Impacted Aggregate	Displays the name of aggregate which is impacted due to the failed disk. By clicking the aggregate name, you can view the aggregate details in the Aggregate details page. If there is no failed disk, no value is displayed in this column.

Storable Capacity	Displays the disk capacity that is available for use.
Raw Capacity	Displays the capacity of the raw, unformatted disk before right-sizing and RAID configuration. By default, this column is hidden.
Type	Displays the types of disks—for example, ATA, SATA, or FCAL.
Effective Type	Displays the disk type assigned by Data ONTAP. Certain Data ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and spare management. Data ONTAP assigns an effective disk type for each disk type.
Firmware	Displays the firmware version of the disk.
RPM	Displays the revolutions per minute (RPM) of the disk. By default, this column is hidden.
Model	Displays the model number of the disk. By default, this column is hidden.
Vendor	Displays the name of the disk vendor. By default, this column is hidden.
Shelf ID	Displays the ID of the shelf where the disk is located. By default, this column is hidden.
Bay	Displays the ID of the bay where the disk is located. By default, this column is hidden.

Related Devices pane

The Related Devices pane enables you to view and navigate to the nodes, SVMs, and aggregates that are related to the cluster:

Nodes	Displays the number, capacity, and health status of the nodes that belong to the selected cluster. Capacity indicates the total usable capacity over available capacity.
Storage Virtual Machines	Displays the number of the SVMs that belong to the selected cluster.
Aggregates	Displays the number, capacity, and the health status of the aggregates that belong to the selected cluster.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts for the selected cluster. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Annotations pane

The Annotations pane enables you to view the annotation for the selected cluster.

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 140
- [Capacity tab](#) on page 140
- [Disk Information tab](#) on page 143
- [Configuration tab](#) on page 143
- [History area](#) on page 144
- [Events list](#) on page 145
- [Related Devices pane](#) on page 145
- [Related Alerts pane](#) on page 145

Command buttons

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

- Actions**
- **Add Alert**
Enables you to add an alert to the selected aggregate.
 - **Edit Thresholds**
Enables you to modify the threshold settings for the selected aggregate.

View Aggregates Enables you to navigate to the Aggregates page.

Capacity tab

The Capacity tab displays detailed information about the selected aggregate, such as its capacity, thresholds, and daily growth rate. By default, the capacity events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to the node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are modified by the technical support representative, the threshold values are applied to the node root aggregate.

- Capacity** Displays the data capacity graph and the Snapshot copies graph, which display capacity details about the aggregate.
- **Snapshot Overflow**
Displays the data space that is consumed by the Snapshot copies.
 - **Used**

Displays the space used by data in the aggregate.

- **Overcommitted**
Indicates that the space in the aggregate is overcommitted.
- **Warning**
Indicates that the space in the aggregate is nearly full. If this threshold is breached, the Space Nearly Full event is generated.
- **Error**
Indicates that the space in the aggregate is full. If this threshold is breached, the Space Full event is generated.
- **Data graph**
Displays the total data capacity and the used data capacity of the aggregate. If the aggregate is overcommitted, a flag is displayed with the overcommitted capacity.
- **Snapshot Copies graph**
This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

Total Capacity	Displays the total capacity in the aggregate.
Data Capacity	Displays the amount of space used by the aggregate (used capacity) and the amount of available space in the aggregate (free capacity).
Snapshot Reserve	Displays the used and free Snapshot capacity of the aggregate.
Overcommitted Capacity	<p>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.</p> <p>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.</p>
Total Cache Space	<p>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.</p> <p>Note: This field is hidden if Flash Pool is disabled for an aggregate.</p>
Aggregate Thresholds	<p>Displays the following aggregate capacity thresholds.</p> <ul style="list-style-type: none"> • Nearly Full Threshold Specifies the percentage at which an aggregate is nearly full.

- Full Threshold
Specifies the percentage at which an aggregate is full.
- Nearly Overcommitted Threshold
Specifies the percentage at which an aggregate is nearly overcommitted.
- Overcommitted Threshold
Specifies the percentage at which an aggregate is overcommitted.

Daily Growth Rate

Displays the disk space used in the aggregate if the rate of change between the last two samples continues for 24 hours.

For example, if an aggregate uses 10 GB of disk space at 2 pm and 12 GB at 6 pm, the daily growth rate (GB) for this aggregate is 2 GB.

Volume Move

Displays the number of volume move operations that are currently in progress.

- Volumes Out
Displays the number and capacity of the volumes that are being moved out of the aggregate.
You can click the link to view more details such as the volume name, aggregate to which the volume is moved, status of the volume move operation, and the estimated end time.
- Volumes In
Displays the number and remaining capacity of the volumes that are being moved into the aggregate.
You can click the link to view more details such as the volume name, aggregate from which the volume is moved, status of the volume move operation, and the estimated end time.
- Estimated used capacity after volume move
Displays the estimated amount of used space (in percentage, and KB, MB, GB, and so on) in the aggregate after the volume move operations are complete.

Capacity Overview - Volumes

Displays graphs that provide information about the capacity of the volumes contained in the aggregate. The amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume is displayed. When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

You can select the graph you want to view from the drop-down lists. You can sort the data displayed in the graph to display details such as the used size, provisioned size, available capacity, the fastest daily growth rate, and the slowest growth rate. You can filter the data based on the Storage Virtual

Machines (SVMs) that contain the volumes in the aggregate. You can also view details for thinly provisioned volumes. You can view the details of specific points on the graph by positioning your cursor over the area of interest. By default, the graph displays the top 30 filtered volumes in the aggregate.

Disk Information tab

The Disk Information tab displays detailed information about the disks in the selected aggregate, including the RAID type and size, and the type of disks used in the aggregate. The tab also graphically displays the RAID groups, the types of disks used (such as SAS, ATA, or FCAL), and the empty slots of the disks that can be added to the aggregate. You can view more information, such as the disk's bay, shelf, and rotational speed, by positioning your cursor over the parity disks and data disks:

- | | |
|---------------------|--|
| RAID Details | <ul style="list-style-type: none"> • Type
Displays the RAID type (RAID0, RAID4, RAID-DP, or Mixed RAID). • Group Size
Displays the maximum number of disks allowed in the RAID group. • Groups
Displays the number of RAID groups in the aggregate. |
| Disks Used | <ul style="list-style-type: none"> • Effective Type
Displays the types of disks (for example, ATA, SATA, or FCAL) in the aggregate. • Data Disks
Displays the number and capacity of the data disks that are assigned to an aggregate. • Parity Disks
Displays the number and capacity of the parity disks that are assigned to an aggregate. |
| Spare Disks | <p>Displays the disk effective type, number, and capacity of the spare disks that are available for the node in the selected aggregate.</p> <p>Note: When an aggregate is failed over to the partner node, Unified Manager does not display all the spare disks compatible with the aggregate.</p> |

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 | <ul style="list-style-type: none"> • Node |
|-----------------|--|

- Displays the name of the node that contains the selected aggregate.
- Block Type
Displays the block format of the aggregate, either 32-bit or 64-bit.
- RAID Type
Displays the RAID type (RAID0, RAID4, RAID-DP, or Mixed RAID).
- RAID Size
Displays the size of the RAID group.
- RAID Groups
Displays the number of RAID groups in the aggregate.
- Flash Pool
Indicates whether or not the aggregate is a Flash Pool.

History area

The History area displays graphs that provide information about the capacity of the selected aggregate.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends—for example, if the aggregate usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Aggregate Capacity Used (%)	Displays the used capacity in the aggregate and the trend in how aggregate capacity is used based on the usage history as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Capacity Used legend, the Capacity Used graph line is hidden.
Aggregate Capacity Used vs Total Capacity	Displays the trend in how aggregate capacity is used based on the usage history, as well as the used capacity, and the total capacity as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Aggregate Capacity Used (%) vs Committed (%) Displays the trend in how aggregate capacity is used based on the usage history, as well as the committed space as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Space Committed legend, the Space Committed graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity Displays the severity of the event.

Event Displays the event name.

Triggered Time Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp when the event was generated is displayed.

Related Devices pane

The Related Devices pane enables you to view the cluster node, volumes, and disks that are related to the aggregate:

Node Displays the capacity and the health status of the node which contains the aggregate. Capacity indicates the total usable capacity over available capacity.

Aggregates in the Node Displays the number and capacity of all the aggregates in the cluster node that contains the selected aggregate. The health status of the aggregates is also displayed, based on the highest severity level. For example, If a cluster node contains ten aggregates, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.

Volumes Displays the number and capacity of the volumes in the selected aggregate. The health status of the volumes is also displayed, based on the highest severity level.

Resource Pool Displays the resource pools related to the aggregate.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected aggregate. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Related tasks

[Performing suggested remedial actions for a full volume](#) on page 69

Job details page

The Job details page enables you to view status and other information about specific protection job tasks that are running, that are queued, or that have completed. You can use this information to monitor protection job progress and to troubleshoot job failures.

Job summary

The job summary displays the following information:

- Job ID
- Type
- State
- Submitted Time
- Completed Time
- Duration

Command buttons

The command buttons enable you to perform the following tasks:

- Refresh** Refreshes the task list and the properties associated with each task.
- View Jobs** Returns you to the Jobs page.

Job tasks list

The Job tasks list displays in a table all the tasks associated with a specific job and the properties related to each task.

- Started Time** Displays the day and time the task started. By default, the most recent tasks are displayed at the top of the column and older tasks are displayed at the bottom.
- Type** Displays the type of task.
- State** The state of a particular task:
 - Completed** The task has finished.
 - Queued** The task is about to run.
 - Running** The task is running.
 - Waiting** A job has been submitted and some associated tasks are waiting to be queued and executed.
- Status** Displays the task status:

- Error** () The task failed.
- Normal** () The task succeeded.
- Skipped** () A task failed, resulting in subsequent tasks being skipped.

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

Definitions of user roles in Unified Manager

The maintenance user or OnCommand administrator assigns a role to every user. Each role contains certain privileges. The scope of activities that you can perform in Unified Manager depends on the role you are assigned and which privileges the role contains.

The following predefined roles exist in Unified Manager:

Operator	Views storage system information and other data collected by Unified Manager, including histories and capacity trends. The role enables the storage operator to view, assign, acknowledge, resolve, and add notes for the events.
Storage Administrator	Configures storage management operations within Unified Manager. The role enables the storage administrator to configure thresholds, create alerts and other storage management-specific options and policies.
OnCommand Administrator	Configures settings unrelated to storage management. The role enables the management of users, security certificates, database access, and administrative options, including authentication, SMTP, networking, and AutoSupport.
Event Publisher	Transmits events generated by partner applications for display in the Unified Manager web UI. This specialized, limited-access user role enables partner applications to share event information with Unified Manager. At the same time, the limited access of this role prevents unauthorized access to the Unified Manager server or the Unified Manager web UI through event publishing activity.

Definitions of user types

A user type specifies the kind of account the user holds and includes remote users, remote groups, local users, database users, and maintenance users. Each of these types has its own role, which is assigned by a user with the role of OnCommand Administrator.

Unified Manager user types are as follows:

Maintenance user	Created from the maintenance console during the initial configuration of Unified Manager. The maintenance user then creates additional users and assigns roles. The maintenance user is also the only user with access to the maintenance console.
Local user	Accesses the Unified Manager UI and performs functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.
Remote group	Groups of users that access the Unified Manager UI using the credentials stored on the authentication server. The name of this account should match the name of a group stored on the authentication server. All users within the remote group are given access to the Unified Manager UI using their individual user credentials. Remote groups can perform functions according to their assigned roles.
Remote user	Accesses the Unified Manager UI using the credentials stored on the authentication server. A remote user performs functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.
Database user	Has read-only access to data in the Unified Manager database, has no access to the Unified Manager web interface or the maintenance console, and cannot execute API calls.

Unified Manager roles and capabilities

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

The following table displays the functions that each role can perform:

Function	Event Publisher	Operator	Storage Administrator	OnCommand Administrator
View storage system information		•	•	•
View other data like histories, capacity trends, and so on		•	•	•

Function	Event Publisher	Operator	Storage Administrator	OnCommand Administrator
View, assign, resolve events		•	•	•
View storage service objects, such as SVM associations and resource pools		•	•	•
Manage storage service objects, such as SVM associations and resource pools			•	•
Define alerts			•	•
Manage storage management options			•	•
Manage storage management policies			•	•
Manage users				•
Manage administrative options				•
Manage database access				•
Publish events	•			

Related references

Definitions of user types on page 148

Definitions of user roles in Unified Manager on page 147

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 150

Diagnostic user capabilities on page 151

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
- Configure network settings
- Change the maintenance user password

Related tasks

Using the maintenance console on page 150

What the maintenance user does

Created during initial configuration, the maintenance user can create subsequent users and assign them roles. The maintenance user can also access the maintenance console and has the role of OnCommand administrator in the web UI.

The maintenance user can perform the following functions using the maintenance console:

- Configure network access
- Upgrade to newer versions of Unified Manager
- Shut down virtual appliances (only from VMware console)

- Increase data disk or swap disk size
- Change the time zone
- Send on-demand AutoSupport messages to technical support from the maintenance console
- Send periodic AutoSupport messages to technical support from the web UI
- Generate support bundles to send to technical support

Related tasks

[Using the maintenance console](#) on page 150

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 150

Accessing the maintenance console using Secure Shell

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

Before you begin

You must have installed and configured Unified Manager.

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

Steps

1. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance.
2. Log in to the maintenance console using your maintenance user name and password.
After 15 minutes of inactivity, the maintenance console logs you out.

Related tasks

[Using the maintenance console](#) on page 150

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 150

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.

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

System Configuration menu

The System Configuration menu enables you to manage your virtual appliance, 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:

Display Server Status	Displays the current server status. Status options include Running and Not Running. If the server is not running, you might need to contact support.
Reboot Virtual Machine	Reboots the virtual machine, stopping all services. After rebooting, the virtual machine and services restart.
Shut Down Virtual Machine	Shuts down the virtual machine, stopping all services. The virtual machine does not restart. You can only select this option from the virtual machine console.
Change <logged in user> User Password	Enables you to change the password of the user currently logged in, which can only be the maintenance user.
Increase Data Disk Size	Enables you to increase the size of your data disks in the virtual machine.
Increase Swap Disk Size	Enables you to increase the size of your swap disks in the virtual machine.
Change Time Zone	Enables you to change the time zone to your location.
Change NTP Server	Enables you to change the NTP Server settings such as IP address or FQDN.
emu	Enables the migration tool user to migrate monitoring data from OnCommand Unified Manager 5.x to OnCommand Unified Manager 6.x. By default, this option is hidden.
Disable Migration Tool user	Disables the migration tool user. This option is displayed only if the migration tool user is enabled.
Back	Returns you to the Main Menu.
Exit	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:

AutoSupport Submission	Enables you to request that an AutoSupport message be generated and sent to technical support or other email recipients.
Post to NetApp	Posts the AutoSupport message to the AutoSupport web server.
Send as email	Sends the AutoSupport message via email to users who need a copy of the data, or to technical support.
Both	Sends the AutoSupport message to both the AutoSupport web server and via email to one or more recipients.
Back	Takes you back to the Main Menu.
Exit	Exits the maintenance console menu.
Generate Support Bundle	Enables you to create a 7-Zip file containing full diagnostic information in the diagnostic user's home directory. The file includes information generated by an AutoSupport message, the contents of the 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 152

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.

Incorrect trigger condition displayed for Aggregate Snapshot Reserve Full event

Issue If the Aggregate Snapshot Reserve Full event is generated in Unified Manager 6.0 and then you upgrade to Unified Manager 6.1, an incorrect trigger condition is displayed in the Summary pane of the Event details page.

Cause There is a mismatch in the number of parameters between Unified Manager 6.0 and Unified Manager 6.1.

Corrective Action Move the event to the resolved state by performing the following steps:

1. Click the **Events** tab.
2. In the Events page, click the event name.
3. In the Event details page, click **Mark As Resolved**.

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.

Issue with installing or regenerating an HTTPS certificate on Unified Manager server enabled for performance monitoring

Issue	If performance monitoring is enabled on the Unified Manager server through connections with one or more Performance Manager servers, and if an HTTPS certificate is installed or regenerated on that Unified Manager server and that server is rebooted, posting of additional performance incidents encountered by Performance Manager servers to the Unified Manager Web UI is stopped.
Cause	The HTTPS certificate installation or regeneration on the Unified Manager server has invalidated the credentials that allowed the Performance Manager servers to post performance incidents to the Unified Manager server.
Corrective actions	You need to delete and then reconfigure the connections between the Unified Manager server and each Performance Manager server that was posting performance incidents to it before the HTTPS certificate installation or regeneration.

Before starting corrective actions, be prepared to respecify the current connection information:

- Unified Manager server name or IP address
- Unified Manager server port (always 443)
- Event Publisher user name (the name of the local Unified Manager server user assigned Event Publisher role privileges)
- Event Publisher password (the password of the local Unified Manager server user assigned Event Publisher role privileges)

Log in to the maintenance console of each Performance Manager server that was posting performance incidents to the Unified Manager server, and then complete the following actions:

1. Type the number of the menu option labeled "Unified Manager Connection" to display the Unified Manager Connection Menu, and then type the number of the menu option labeled "Delete Unified Manager Server Connection."
2. When prompted to confirm that you want to continue the deletion, type **y**, and then press any key to continue.
3. Type the number of the menu option labeled "Add / Modify Unified Manager Server Connection."
4. When prompted, supply the requested Unified Manager server name or IP address and Unified Manager server port information.
5. When prompted, accept the Unified Manager server trust certificate to support the connection. Do not change the default port value (443).

6. When prompted, supply the requested Event Publisher user name and Event Publisher password, and then confirm that the settings are correct.
7. To exit the maintenance console, press any key to continue, and then type **x**.

Glossary

A

Access Control List (ACL)	A set of data associated with a file, directory, or other resource or share that defines user or group access rights to that resource or share.
admin SVM	Formerly known as admin Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that has overall administrative access to all objects in the cluster, including all objects owned by other SVMs, but does not provide data access to clients or hosts.
aggregate	A set of multiple RAID (Redundant Array of Inexpensive Disks or Redundant Array of Independent Disks) groups that can be managed as a single unit for protection and provisioning purposes.
aggregate committed capacity	The data storage space in an aggregate that is committed to provide for its underlying volumes. Calculated by the total capacity provisioned for volumes.
aggregate total capacity	The data storage space within an aggregate that can be used by volumes or aggregate-level Snapshot copies. Calculated by the total data capacity of the aggregate plus the aggregate-level Snapshot reserve space.
alert	<ul style="list-style-type: none"> • In OnCommand Insight (formerly SANscreen suite), an alarm indicating that a device or path state has changed in a way that violates a policy or exceeds a threshold. • In Unified Manager, a user-configured notification that is sent whenever a specific event or an event of a specific severity type occurs, not necessarily related to a specific user. Alerts are used to monitor and manage datasets and resources. See also <i>event</i> and <i>severity type</i>.
AutoSupport	An integrated technology that triggers email messages from the customer site to technical support or another specified email recipient when there are any failures in Unified Manager services. These messages contain information such as feature usage metrics, configuration and user settings, system health, and so on.
available capacity	The amount of usable space available in a storage system. Calculated by the used capacity minus the unused reserve capacity.

B

- backup relationship** A persistent association between a primary directory and a secondary volume for disk-based data backup and restore using the Data ONTAP SnapVault feature.
- baseline transfer** An entire transfer of data as compared to an incremental transfer of data.

C

- CIFS** See *Common Internet File System (CIFS)*.
- CIFS share**
- In Data ONTAP, a directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known simply as a *share*.
 - In OnCommand Insight (formerly SANscreen suite), a service exposed from a NAS device to provide file-based storage through the CIFS protocol. CIFS is mostly used for Microsoft Windows clients, but many other operating systems can access CIFS shares as well.
- client application** An application that calls Unified Manager APIs to enable its operator to configure, monitor, and initiate data management operations to be executed on the Unified Manager server.
- cluster**
- In clustered Data ONTAP 8.x, a group of connected nodes (storage systems) that share a namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits.
 - In the Data ONTAP 7.1 release family and earlier releases, a pair of storage systems (sometimes called *nodes*) configured to serve data for each other if one of the two systems stops functioning.
 - In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.
 - For some storage array vendors, *cluster* refers to the hardware component on which host adapters and ports are located. Some storage array vendors refer to this component as a *controller*.
- cluster committed capacity** The data storage space in a cluster that is committed to provide for its underlying aggregates. Calculated by the sum of the total capacity of all the aggregates in the cluster.
- cluster failover (CFO)** In Data ONTAP 7.1.x and earlier, the method of ensuring data availability by transferring the data service of a failed node to another node in an . 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 are connected, and over which heartbeat and WAFL log information are transmitted when both nodes are running.
cluster total capacity	The data storage space in a cluster that can be used by aggregates or volumes. Calculated by the sum of the capacity of all the data disks excluding disk right-sizing and reservation plus sum of the capacity of all spare disks excluding right-sizing.
cluster Vserver	Former name for a data SVM; see data SVM.
container object	An object, such as an aggregate or a Storage Virtual Machine (SVM, formerly known as Vserver), in which data objects reside.
counter	The statistical measurement of activity on a storage system or storage subsystem that is provided by Data ONTAP. Each type of storage system or subsystem has a set of counters.

D

data capacity	The storage space that is set aside by the container, such as an aggregate or a volume, to store user data. Typically, this capacity can be used for any container, but data is only written to the lowest level container, usually the volume.
Data ONTAP	The operating system software running on NetApp storage devices.
datastore	A storage location for virtual machines, such as a VMFS volume, a directory on a NAS server, or a local file system path. A datastore is platform-independent and host-independent; therefore, it does not change when a virtual machine it contains moves to another host.
data object	A container of data, such as a file, directory, volume, or LUN, that can be discovered, monitored, protected, created, or restored by the Unified Manager server.
data SVM	Formerly known as data Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data SVMs within a cluster.
deduplication	The consolidation of blocks of duplicate data into single blocks to store more information using less storage space.

deduplication return	The capacity savings resulting from deduplication. Calculated by the volume capacity before deduplication - the volume capacity after deduplication.
destination	The storage to which source data is backed up, mirrored, or migrated.
destination data object	A data object that contains the backed up or mirrored replicated data.
dedupe	See <i>deduplication</i> .
DHCP	See <i>Dynamic Host Configuration Protocol (DHCP)</i> .
Dynamic Host Configuration Protocol (DHCP)	The protocol for automating the assignment of network addresses.

E

ESX server	A VMware term describing a server that abstracts server processor, memory, storage, and networking resources into multiple virtual machines.
event	An indication of a predefined condition occurring or when an object crosses a threshold. All events are assigned a severity type and are automatically logged in the Events window. See also <i>alert</i> and <i>severity type</i> .

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 <i>Fully Qualified Domain Name (FQDN)</i> .
Fully Qualified Domain Name (FQDN)	The complete name of a specific computer on the Internet, consisting of the computer's host name and its domain name.
fractional reserve	An option that determines how much space in a volume is reserved for Snapshot overwrite data for LUNs and space-reserved files, to be used after all other space in the volume is used.

G

- giveback** The return of identity from an emulated storage system to the failed system, resulting in a return to normal operation. The reverse of *takeover*.
- global namespace** See *namespace*.
- growth rate** The measurement of how fast the storage is filling. The growth rate is determined by dividing the daily growth rate by the total amount of space in the storage system.

H

- HA (high availability)**
- In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an *HA pair*, that are configured to serve data for each other if one of the two nodes stops functioning.
 - In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.
- HA pair**
- In Data ONTAP 8.x, a pair of nodes whose controllers are configured to serve data for each other if one of the two nodes stops functioning. Depending on the system model, both controllers can be in a single chassis, or one controller can be in one chassis and the other controller can be in a separate chassis.
 - In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.
- host** A computer system that accesses data on a storage system.
- host bus adapter (HBA)** An interface card that plugs into a SAN device. SAN devices use the ports on their respective HBAs to connect to each other in a SAN. Each SAN device might contain one or more HBAs, and an HBA might contain more than one port. Each port can be used to establish a connection to a SAN.

I

- igroup** initiator group. A collection of unique iSCSI node names of initiators (hosts) in an IP network that are given access to *front-end LUNs* when they are mapped to those LUNs. (Array LUNs on a storage array that provide storage for V-Series systems can be considered *back-end LUNs*.)

incident	An issue that has already impacted the availability or capacity of storage objects.
incremental transfer	A subsequent backup after a baseline transfer has occurred of a primary directory in which only the new and changed data since the last backup (baseline or incremental) is transferred. The transfer time of incremental transfers can be significantly less than the baseline transfer.
initiator	The system component that originates an I/O command over an I/O bus or network. The target is the component that receives this command.
inode	A data structure containing information about files on a storage system and in a UNIX file system.
iSCSI	Internet Small Computer Systems Interface (iSCSI) protocol. A licensed service on the storage system that enables you to export LUNs to hosts using the SCSI protocol over TCP/IP.
iSCSI router	A storage router implementing the Internet Small Computer Systems Interface (iSCSI) protocol (SCSI over IP) to extend access of a Fibre Channel fabric and attached storage devices to IP servers.

J

JBOD	Just a Bunch Of Disks. An array of disks without any redundancy; that is, without RAID configuration.
job	A long-running operation, for example, scheduled local backup of a dataset, a mirror transfer, and password updates.

L

level-0 backup	An initial backup (also known as a <i>baseline transfer</i>) 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 <i>VIF</i> (virtual interface) in Data ONTAP GX. A logical network interface, representing a network access point to a node. LIFs currently correspond to IP addresses, but could be implemented by any interconnect. A LIF is generally bound to a physical network port; that is, an Ethernet port. LIFs can fail over to other physical ports (potentially on other nodes) based on policies interpreted by the LIF manager.
Lightweight Directory Access Protocol (LDAP)	A client-server protocol for accessing a directory service.

local backup	Local backup protection (also referred to as <i>Snapshot protection</i>) 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 <i>virtual disks</i> , and vice versa. See also <i>virtual disk</i> .
logical object	The entity that represents a storage container, such as a volume, qtree, LUN, or dataset.
lower threshold	The value set to generate an event when a counter falls and remains below that value for longer than the specified interval.

M

A (v) next to a term indicates that the definition is for the verb form of the word, while an (n) next to a term indicates that the definition is for the noun form of the word.

maintenance user	The user who has access rights to deploy and configure an OnCommand Unified Manager virtual appliance.
Management Information Base (MIB)	ASCII files that describe the information that the SNMP agent sends to network management stations.
member	Any data object that subscribes to or is created by a storage service.
mirror (v)	The process of creating an exact duplicate of all volume data from a NetApp source storage system to a destination storage system. If data in the source storage system is lost or made unavailable, then that replicated data can quickly be made available from the destination mirror site. Mirror operations are employed from primary to secondary storage and from secondary to tertiary storage. Mirror protection is based on NetApp Volume SnapMirror technology.
mirror copy (n)	The exact duplicate of all volume data (both active and protected) from a NetApp source storage system to a destination storage system, created using NetApp Volume SnapMirror technology.

move (v) To physically move data and any needed associated configuration of an object from one aggregate to another within a cluster, including within a single node.

N

namespace In network-attached storage (NAS) cluster environments, an abstraction layer for data location that provides a single access point for all data in the system. It enables users to access data without specifying the physical location of the data, and enables administrators to manage distributed data storage as a single file system. Sometimes referred to as *global namespace*.

NDMP Network Data Management Protocol. A protocol that allows storage systems to communicate with backup applications and provides capabilities for controlling the robotics of multiple tape backup devices.

Network File System (NFS) export A service exposed from a NAS device to provide file-based storage through the NFS protocol. NFS is mostly used for UNIX-like operating systems, but other operating systems can access NFS exports as well.

node In Data ONTAP, a storage system or in a cluster or an HA pair. To distinguish between the two nodes in an HA pair, one node is sometimes called the *local node* and the other node is sometimes called the *partner node* or *remote node*.

nondisruptive The ability of a system to continue serving data to clients during a system process or activity, such as a LUN restore operation or an online migration.

O

offline A database state indicating that the database is not available to users (for example, the database is in the following states: shutdown, started, or mounted).

online A database state indicating that the database is available to users (for example, open).

OnCommand administrator An RBAC role that enables a person to configure settings for items unrelated to storage management, such as user roles, security certificates, database access, LDAP, SMTP, networking, and AutoSupport.

operator An RBAC role that enables a person to view data and to view, assign, and resolve events in OnCommand Unified Manager.

P

parity disk	The disk on which parity information is stored for a RAID4 disk drive array. In RAID groups using RAID-DP protection, two parity disks store the parity and double-parity information. Used to reconstruct data in failed disk blocks or on a failed disk.
partner node	From the point of view of the local node (storage system), the other node in .
port	A physical connection point on computers, switches, storage arrays, and so on, which is used to connect to other devices on a network. Ports on a Fibre Channel network are identified by their World Wide Port Name (WWPN) IDs. TCP/IP ports are used as virtual addresses assigned to each IP address.
policy	A set of parameters that are grouped together as a distinct entity, so that the set of parameters can be applied to objects as a unit.
protection artifact	An object, such as a destination data object, or a protection relationship that the Unified Manager server creates to support protection jobs when a data object is subscribed to a storage service.
protection policies	The entities that enable you to set the automation controls for scheduling, monitoring, and alerts on any set of data in terms of normal backup, offsite backup, disaster and recovery backup, and regulatory copies.
protection relationship	The SnapMirror or SnapVault relationship that exists between a source data object and a destination data object.

Q

qtree	A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.
--------------	--

R

RAID-DP	redundant array of independent disks, double-parity.
raw capacity	The total amount of addressable blocks on physical disk drives. Calculated by multiplying the number of disk drives by the labeled capacity of those disk drives. For V-Series systems, it refers to the size of the array LUNs.
RBAC	Role-based Access Control. A system whereby access to resources is decided based on the role of a user. RBAC controls who has access to various

	operations on which resources. Access to resources is first assigned to roles and roles are then assigned to users. Conforms to NIST RBAC standard.
recovery	The re-creation of a past operational state of an entire application or computing environment. Compare <i>restore</i> . A recovery operation can encompass a restore. <i>Recovery</i> and <i>restore</i> are often used synonymously. Recovery and restore can also be <i>in-place</i> , 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 <i>secondary storage</i> . See also <i>local backup</i> .
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 <i>clones</i> , <i>point-in-time copies</i> , or <i>Snapshot copies</i> , 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 <i>recovery</i> . A restore can be part of a recovery operation. <i>Restore</i> and <i>recovery</i> are often used synonymously. Restore and recovery can also be <i>in-place</i> , meaning that copies are mounted and used locally instead of being copied elsewhere.
retention period	The user-specified minimum length of time that a local backup Snapshot copy must be retained.
risk	The issues that can impact the availability or capacity of storage objects.
root member	A data object that subscribes to a storage service.

S

SAN	storage area network. A dedicated network linking servers or workstations to devices, typically over Fibre Channel. SAN allows data and devices to be shared across a network as if they were attached locally.
severity type	The level of priority assigned to an event to help determine priorities for taking corrective action.
SFO	See <i>storage failover (SFO)</i> .
share	A directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known as a <i>CIFS share</i> .
Snapshot available capacity	The storage space that is available in the Snapshot reserve for Snapshot copies. Calculated by subtracting the total Snapshot used capacity from the Snapshot reserve capacity.

Snapshot copy	An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copies enable users to restore files and to back up the storage system to tape while the storage system is in use.
Snapshot overflow	The storage space that is consumed by Snapshot copies from the total data capacity of a volume or an aggregate. Calculated by subtracting the Snapshot reserve capacity from the Snapshot used capacity.
Snapshot reserve capacity	The storage space that is set aside by the volume or the aggregate for its Snapshot copies. Data cannot be written to this space.
Snapshot return	The capacity savings of Snapshot copies when compared to full volume copies. Calculated by the volume capacity - the Snapshot capacity.
Snapshot unused capacity	The Snapshot reserve space remaining after the Snapshot copies are created.
Snapshot used capacity	The storage space used by the Snapshot copies in a volume or an aggregate.
spare disk	A physical disk that is part of a storage device that is the same technology type (FC, SATA), size, and speed as a standard disk. A spare disk is used in case the standard disk malfunctions.
storable capacity	The disk capacity that is available for use after right-sizing. A configuration set up so that one node automatically takes over for its partner when the partner node becomes impaired.
storage administrator	Configures storage management operations within Unified Manager. The role enables the storage administrator to configure thresholds and to create alerts and other storage management-specific options and policies.
storage controller	The component of a storage system that runs the operating system and controls its disk subsystem. Storage controllers are also sometimes called <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> .
storage efficiency	The ratio of usable capacity to effective used capacity, accounting for efficiency returns. Calculated by the effective used capacity / the usable capacity.
storage failover (SFO)	The method of ensuring data availability by transferring the data service of a failed node to another node in the cluster. Transfer of data service is often transparent to users and applications. Also referred to as <i>controller failover (CFO)</i> or <i>cluster failover (CFO)</i> .
storage utilization	The ratio of usable capacity to used capacity, without accounting for efficiency returns. Calculated by the used capacity / the usable capacity.

Storage Virtual Machine (SVM)	(Known as <i>Vserver</i> prior to clustered Data ONTAP 8.2.1. The term “Vserver” is still used in CLI displays and <code>vserver</code> command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs— <i>admin</i> , <i>node</i> , and <i>data</i> —but unless there is a specific need to identify the type of SVM, “SVM” usually refers to the data SVM.
SVM	(Storage Virtual Machine; known as <i>Vserver</i> prior to clustered Data ONTAP 8.2.1. The term “Vserver” is still used in CLI displays and <code>vserver</code> command syntax.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of SVMs— <i>admin</i> , <i>node</i> , and <i>data</i> —but unless there is a specific need to identify the type of SVM, “SVM” usually refers to the data SVM.
SVM guaranteed available capacity	The data storage space that is guaranteed by the SVM to its underlying volumes but is not used. Calculated as the sum of the available size of all the thick provisioned volumes.
SVM used capacity	The data storage space that is guaranteed by the SVM to its underlying volumes. Calculated as the sum of the used data capacity of all the volumes associated with the SVM.
SVM unguaranteed capacity	The data storage space that is not guaranteed by the SVM to its underlying volumes. Calculated as the sum of the available sizes of all the thin provisioned volumes.
SVM total capacity	The sum of the total data storage space of all the volumes in the SVM.
SVM unguaranteed capacity	The data storage space that is not guaranteed by the SVM to its underlying volumes. Calculated as the SVM total capacity minus the sum of the committed capacity of aggregates that are associated with the SVM.
system reserve capacity	The capacity required for fixed system reserves, RAID parity, mirroring, and spare drives. Calculated by the fixed reserve + RAID reserve + spares.

T

takeover	The emulation of the failed node identity by the takeover node in ; the opposite of <i>giveback</i> .
takeover node	A node (storage system) that remains in operation after the other node stops working and that hosts a virtual node that manages access to the failed node disk shelves and network connections. The takeover node maintains its own identity and the virtual node maintains the failed node identity.

thin provisioning (TP) A method of optimizing the efficiency with which the available space is used in storage. When many applications share access to the same storage array, thin provisioning enables administrators to maintain a single free space buffer pool to service the data requirements of all applications. This is done by allocating disk storage space in a flexible manner among multiple consumers, based on the minimum space required by each at any given time. This avoids poor utilization that occurs on traditional storage arrays where large pools of storage capacity are allocated to individual applications, but much remains unused.

trap An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the storage system.

U

unassigned disk A disk that is not assigned to any node or counted as a spare disk.

unused capacity The free, usable storage space available for storing user data on the device, excluding capacity reserved for Snapshot copies or data.

unused reserve capacity The capacity allocated but unused by Aggregate Snapshot Reserve, Volume Snapshot Reserve, Volume Fractional Reserve, and Vol/LUN/File Guaranteed Space. These reserves are adjustable by the user. Calculated by the aggregate snapshot unused reserve + volume snapshot unused reserve + volume fractional unused reserve + vol/LUN/file unused guaranteed space.

usable capacity The capacity available to applications and users. Calculated by the raw capacity - the system reserve capacity.

used capacity The capacity used by application or user data, including volume Snapshot copies and aggregate Snapshot copies. Calculated by the usable capacity - the free capacity.

V

vApp See *virtual appliance*.

virtual appliance A prebuilt software solution containing virtual machines and software applications that are integrated, managed, and updated as a package. Also called *vApp*.

virtual machine A guest operating system and any application installed thereon, running on a computing device on which the software is installed, or suspended to disk or any other storage media accessible by the computing device.

**VMware
VirtualCenter
(VC)**

A management software suite used to create your VMware datastores and virtual machines and to configure the storage system volumes as the containers in which your active datastore and virtual machine images reside. It consists of the following components:

- VMware agents—software modules installed on an ESX server to carry out VC server requests
- VirtualCenter (VC) server—a server communicating with VMware agents on an ESX server
- Virtual Infrastructure (VI) client—a GUI client to manage the VC server
- VMware ESX server—an enterprise-level product that integrates server processes, storage functionality, and networking resources into multiple virtual systems.

ESX server can also refer to a physical host running an ESX server hypervisor OS.

volume

- For Data ONTAP, a logical entity that holds user data that is accessible through one or more of the supported access protocols, including Network File System (NFS), Common Internet File System (CIFS), Fibre Channel (FC), and Internet SCSI (iSCSI). V-Series treats an IBM volume as a disk.
- For IBM, the area on the storage array that is available for a V-Series system or non V-Series host to read data from or write data to. The V-Series documentation uses the term *array LUN* to describe this area.

**volume committed
capacity**

The data storage space in a volume that is committed to provide storage space for its underlying qtrees based on their quota settings. Calculated as the sum of all qtrees disk hard limits. The sum does not include the qtrees for which the disk hard limit is not set.

**volume total
capacity**

The data storage space in a volume that can be used by qtrees, LUNs, or other files and volume-level Snapshot copies. Calculated as the total data capacity of the volume plus the volume-level Snapshot reserve space.

Vserver

(Known as “Storage Virtual Machine (SVM)” in clustered Data ONTAP 8.2.1 and later.) A virtual machine that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of Vservers—*admin*, *node*, and *cluster* (“cluster Vserver” is called “data Vserver” in Data ONTAP 8.2)—but unless there is a specific need to identify the type of Vserver, “Vserver” usually refers to the cluster/data Vserver.

Copyright information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bypass, 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 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

6.1 release of Unified Manager
related concepts [8](#)

- A**
- acknowledging
 - events [97](#)
 - adding
 - alerts [38, 63](#)
 - authentication servers [34](#)
 - clusters [40, 59](#)
 - new rules [66](#)
 - notes about an event [96](#)
 - rules [66](#)
 - administrative tasks
 - summary of common workflows for performing [25](#)
 - administrators
 - OnCommand [147](#)
 - storage [147](#)
 - Aggregate details page [140](#)
 - Aggregate Snapshot Reserve Full event
 - troubleshooting incorrect trigger condition [157](#)
 - aggregates
 - configuring global threshold values for [35](#)
 - details about [140](#)
 - viewing annotations for [90](#)
 - alerts
 - adding [38, 63](#)
 - configuring your environment for [31](#)
 - creating [38, 63](#)
 - alerts not received
 - troubleshooting [158](#)
 - annotations
 - creating rules for storage objects [89](#)
 - definition of [19](#)
 - description of types [91](#)
 - for volumes, details [102](#)
 - removing [90](#)
 - using for storage objects [89](#)
 - viewing for storage objects [90](#)
 - appliances
 - overview of backup and restore process for virtual [18](#)
 - assigning
 - events [96](#)
 - authentication
 - adding servers [34](#)
 - enabling remote [33](#)
 - AutoSupport
 - using the maintenance console [154](#)
 - availability events
 - general description [15](#)
 - availability health
 - defined [15](#)
 - availability issues
 - correcting a flash card offline condition [42, 43](#)
 - correcting a storage failover interconnect link down condition [44, 46](#)
 - resolving volume offline conditions [47, 48](#)
 - availability workflows
 - introduction to [42](#)
- B**
- backup process
 - overview of virtual appliance [18](#)
 - backup vault protection
 - configuration [17](#)
 - backup vault protection relationships
 - defined [16](#)
 - bundles
 - generating support [93](#)
- C**
- capabilities
 - database user [15](#)
 - FlexVol volume [12](#)
 - table of roles associated with [148](#)
 - capacity
 - information for volumes [102](#)
 - capacity events
 - full volume [69](#)
 - resolving [68](#)
 - suggested remedial actions for a full volume [69](#)
 - certificates
 - generating HTTPS security certificates [29](#)
 - viewing HTTPS security [30](#)
 - cluster
 - description of [9](#)
 - cluster components

- definition [21](#)
- limits imposed on workloads by [21](#)
- Cluster details page [131](#)
- clustered Data ONTAP systems
 - See* clusters
- clusters
 - adding [40, 59](#)
 - creating rules to annotate [89](#)
 - details about [131](#)
 - viewing discovery status [40, 59](#)
- configuring
 - aggregate global threshold values [35](#)
 - DNS [28](#)
 - information for volumes [102](#)
 - network settings [28](#)
 - notification settings [32](#)
 - thresholds [35](#)
 - volume global threshold values [36](#)
 - your environment [26](#)
- connections
 - between Performance Manager and Unified Manager, purpose of [23](#)
- creating
 - alerts [38, 63](#)
 - custom rules [67](#)
 - protection relationships [70](#)
 - rules [66](#)
 - rules using templates [66](#)
- custom rules
 - creating [67](#)

D

- data
 - restoring from the Volume details page [87](#)
 - restoring from the Volumes page [88](#)
- data disk size
 - increasing [154](#)
- data policies
 - configuration details [116](#)
 - defined [14](#)
 - exporting [68](#)
 - workflow for managing Infinite Volumes with storage classes [61](#)
- database users
 - capabilities [15](#)
 - creating [37](#)
 - defined [148](#)
- definitions
 - performance incidents [20](#)

- deleting
 - annotations from storage objects [90](#)
- DHCP
 - enabling [28](#)
- diag users [151](#)
- diagnostic information
 - generating support bundles containing full [93](#)
- discovery
 - viewing status of cluster [40, 59](#)
- DNS
 - configuring [28](#)

E

- editing
 - Infinite Volume threshold settings [60](#)
 - network settings [28](#)
 - storage class threshold settings [63](#)
 - unmanaged relationship lag threshold settings [37](#)
- efficiency
 - information for volumes [102](#)
- enabling
 - DHCP [28](#)
- environment
 - setup [26](#)
- error events
 - performing suggested remedial actions for a full volume [69](#)
- Event details page [98](#)
- event impact areas
 - availability [101](#)
 - capacity [101](#)
 - configuration [101](#)
 - description [101](#)
 - performance [101](#)
 - protection [101](#)
- event impact levels
 - description [101](#)
 - event [101](#)
 - incident [101](#)
 - risk [101](#)
- event publishers
 - defined [147](#)
- event severity types
 - critical [101](#)
 - description [101](#)
 - error [101](#)
 - information [101](#)
 - warning [101](#)
- event states

definition of [19](#)

events

acknowledged event, definition [19](#)

acknowledging [97](#)

adding notes about [96](#)

assigning to users [96](#)

definition of [18](#)

details [98](#)

generated by performance incidents [20](#)

impact areas [101](#)

impact levels [101](#)

new event, definition [19](#)

obsolete event, definition [19](#)

performing suggested remedial actions for a full volume [69](#)

resolved event, definition [19](#)

resolving [97](#)

reviewing notes about [96](#)

severity types [101](#)

states defined [19](#)

viewing notes about [96](#)

F

failed protection jobs

identifying [81](#)

identifying the cause [82](#)

performing corrective actions [82](#)

resolving [81](#)

failover and failback

using reverse resync for [76](#)

failure

protection job [81](#)

Filezilla

using to retrieve support bundles [93](#)

flash card offline conditions

troubleshooting [42, 43](#)

FlexVol volumes

capabilities of [12](#)

defined [12](#)

with SVMs, explained [9](#)

H

high annotation types

definition of [91](#)

host names

changing [27](#)

HTTPS

viewing the security certificate [30](#)

HTTPS certificates

generating new security certificates [29](#)

troubleshooting issue with installing or regenerating on a Unified Manager server enabled for performance monitoring [159](#)

I

I/O performance incidents

defined [20](#)

impact areas

availability [101](#)

capacity [101](#)

configuration [101](#)

description [101](#)

performance [101](#)

protection [101](#)

impact levels

description [101](#)

event [101](#)

incident [101](#)

risk [101](#)

incidents

I/O performance, defined [20](#)

Infinite Volumes

creating custom rules for [67](#)

creating rules for placing data in [66](#)

definition of [13](#)

editing threshold settings [60](#)

storage classes, definition of [13](#)

with SVMs, explained [9](#)

workflow for managing [61](#)

workflow for monitoring [58](#)

workflow for setting up [58](#)

issues

summary of workflows for troubleshooting common [25](#)

J

Job Details page

purpose [146](#)

jobs

defined [13](#)

identifying cause of failure [81](#)

list of those you can monitor [13](#)

progress monitoring [146](#)

resolving terminated [81](#)

status [146](#)

troubleshooting failures [146](#)

L

- lag issues
 - resolving [85](#)
- local users
 - changing password for [41](#)
 - creating [37](#)
 - defined [148](#)
- low annotation types
 - definition of [91](#)

M

- maintenance console
 - accessing using Secure Shell [92, 151](#)
 - accessing using VM console [152](#)
 - AutoSupport [154](#)
 - diag users
 - capabilities [151](#)
 - generating a support bundle using [93](#)
 - overview [7](#)
 - purpose [150](#)
 - restarting the virtual machine [31](#)
 - restarting Unified Manager [31](#)
 - role of maintenance user [150](#)
 - support and diagnostics [154](#)
 - system configuration [154](#)
 - using for configuration [152](#)
 - what it does [150](#)
- maintenance user
 - defined [148](#)
 - what this user does [150](#)
- managing
 - Infinite Volumes, workflow for [61](#)
- migration tool
 - enabling [154](#)
- migration tool user
 - disabling [154](#)
- mirror protection
 - configuration [17](#)
- mirror protection relationships
 - defined [16](#)
- mission-critical annotation types
 - definition of [91](#)
- modifying
 - Infinite Volume threshold settings [60](#)
 - storage class threshold settings [63](#)
 - unmanaged relationship lag threshold settings [37](#)
- monitoring
 - Infinite Volumes, workflow for [58](#)

protection relationships [70](#)

N

- Network Configuration menu
 - using the maintenance console [152](#)
- network interfaces
 - adding new [155](#)
- network settings
 - configuring [28](#)
 - customizing the host name [27](#)
 - editing [28](#)
- nodes
 - single node cluster
 - See* single-node cluster
- notification
 - adding alerts [38, 63](#)
 - configuring settings [32](#)

O

- offline flash card conditions
 - troubleshooting [42, 43](#)
- offline volumes
 - determining if caused by a down host cluster node [49](#)
 - determining if caused by a stopped SVM resulting from a down cluster node [50](#)
 - determining if caused by broken RAID disks [52](#)
- OnCommand administrators
 - defined [147](#)
- OnCommand Workflow Automation
 - integrating with OnCommand Unified Manager [71](#)
 - pairing with Unified Manager [71](#)
- operators
 - defined [147](#)

P

- pairing
 - Workflow Automation [71](#)
- passwords
 - changing local user [41](#)
- performance
 - diagnosing issues with [53](#)
- performance incidents
 - analysis workflow [55](#)
 - definition of [20](#)
 - display of, in Unified Manager [20](#)
 - generation on a cluster component [21](#)

Performance Manager

- about connections between multiple Performance Manager servers, and Unified Manager [24](#)
- configuring a connection to a Unified Manager server [54](#)
- purpose of connection with Unified Manager [23](#)
- using to monitor performance [53](#)

performance monitoring

- configuring connections between Performance Manager and Unified Manager [54](#)
- enabling [54](#)
- example workflow [55](#)
- troubleshooting disablement due to HTTPS certificate installation or regeneration [159](#)

physical storage

- adding clusters [40](#), [59](#)

policies

- creating for SnapMirror relationships [75](#)
- creating for SnapVault relationships [74](#)
- exporting data policies [68](#)

protection

- information for volumes [102](#)

protection job failures

- identifying [81](#)
- identifying the cause [82](#)
- performing corrective actions [82](#)
- resolving [81](#)

protection jobs

- correcting failed [82](#)
- execution [17](#)
- identifying failed [81](#)
- resolving failed [81](#)

protection relationships

- creating [70](#), [71](#)
- creating from the Volume details page [72](#), [73](#)
- creation [17](#)
- lag issue resolution workflow [85](#)
- monitoring [70](#)
- removing from the Volume details page [80](#)
- reversing from the Volume details page [78](#)
- troubleshooting [70](#)

purpose of maintenance console

- list of actions performed using [150](#)

R

reassigning

- events [96](#)

reference information

- common to workflows [96](#)

relationships

- unmanaged, editing lag thresholds settings for [37](#)

releases of Unified Manager

- concepts related to working with 6.1 [8](#)

remote authentication

- enabling [33](#)

remote groups

- adding [37](#)
- defined [148](#)

Remote User UI option

- troubleshooting lack of display [157](#)

remote users

- adding [37](#)
- defined [148](#)

resolving

- events [97](#)

resource pools

- about [14](#)

resources

- selecting using SVM associations [15](#)

restore process

- overview of virtual appliance [18](#)

resynchronization operations

- selecting maximum transfer rate for [80](#)
- selecting Snapshot copies for [80](#)
- selecting transfer priority for [80](#)

reverse resync

- using in failover and failback scenarios [76](#)

reverse resync operations

- performing from the Volume details page [78](#)

reviewing

- notes about events [96](#)

role-based access control

- See* RBAC

roles

- assigning to users [37](#)
- creating a user with the Event Publisher user [53](#)
- defined [147](#)
- table of capabilities associated with [148](#)

rules

- adding [66](#)
- creating to annotate storage objects [89](#)
- creating, custom [66](#), [67](#)
- creating, using templates [66](#)
- defined [14](#)
- exporting [68](#)

rules using templates

- creating [66](#)

S

- schedules
 - creating for SnapMirror transfers [76](#)
 - creating for SnapVault transfers [76](#)
- Secure Shell
 - using to access the maintenance console [92](#), [151](#)
- security certificates
 - generating, HTTPS [29](#)
 - viewing HTTPS [30](#)
- servers
 - about connections between multiple Performance Manager servers and Unified Manager [24](#)
- setting up
 - aggregate global threshold values [35](#)
 - notification settings [32](#)
 - SMTP server [32](#)
 - SNMP [32](#)
 - thresholds [35](#)
 - volume global threshold values [36](#)
- setup
 - post-deployment [26](#)
- severity types
 - critical [101](#)
 - description [101](#)
 - error [101](#)
 - information [101](#)
 - warning [101](#)
- single-node cluster
 - description of [9](#)
- SnapMirror license
 - enabling mirror protection [16](#)
- SnapMirror relationships
 - breaking [78](#)
 - breaking before a reverse resync [76](#)
 - creating [71](#)
 - creating from the Volume details page [72](#)
 - creating policies for [75](#)
- SnapShot copies
 - restoring data from [87](#)
 - selecting for resynchronization operations [80](#)
- SnapVault license
 - enabling backup vault protection [16](#)
- SnapVault policies
 - creating [74](#)
- SnapVault relationships
 - creating [71](#), [73](#)
 - creating policies for [74](#)
- storage administrators
 - defined [147](#)
- storage classes
 - capacity details of [116](#)
 - definition of [13](#)
 - editing threshold settings [63](#)
 - workflow for managing Infinite Volumes with [61](#)
- storage failover link down conditions
 - troubleshooting [44](#), [46](#)
- storage object events
 - prioritizing with annotations [89](#)
- storage objects
 - creating rules to annotate [89](#)
 - removing annotations [90](#)
 - viewing annotations for [90](#)
- Storage Virtual Machine details page [116](#)
- support bundles
 - generating [93](#)
 - introduction to sending to technical support [91](#)
 - retrieving using a Windows client [93](#)
 - retrieving using the CLI [94](#)
 - sending to technical support for diagnosis [93](#)
 - uploading to technical support [95](#)
- SVM associations
 - about [15](#)
- SVMs
 - creating rules to annotate [89](#)
 - details about [116](#)
 - viewing annotations for [90](#)
 - with FlexVol volumes, explained [9](#)
 - with Infinite Volume, explained [9](#)
- SVMs with FlexVol volumes
 - explained [9](#)
- SVMs with Infinite Volume
 - explained [9](#)
- swap disk size
 - increasing [154](#)
- System Configuration menu
 - changing user password [154](#)
 - displaying server status [154](#)
 - rebooting virtual machine [154](#)
 - shutting down virtual machine [154](#)
 - using the maintenance console [152](#)

T

- tasks
 - common to workflows [96](#)
 - summary of common workflows for performing administrative [25](#)
 - viewing information about [146](#)
- technical support

- generating support bundles for [93](#)
- introduction to sending support bundles to [91](#)
- thresholds
 - configuring [35](#)
 - editing settings for Infinite Volumes [60](#)
 - editing settings for storage classes [63](#)
 - editing settings for unmanaged relationships [37](#)
 - global values for aggregates [35](#)
 - global values for volumes [36](#)
- time zone
 - changing [154](#)
- tools for retrieving support bundles
 - Filezilla [93](#)
 - WinSCP [93](#)
- transfer priorities
 - specifying for SnapMirror relationships [75](#)
- transfers
 - creating schedules for data protection [76](#)
- troubleshooting
 - alerts not received by designated recipients [158](#)
 - flash card offline condition [43](#)
 - flash card offline conditions [42](#)
 - generating support bundles for technical support [93](#)
 - incorrect trigger condition for Aggregate Snapshot Reserve Full event [157](#)
 - introduction to sending support bundles to technical support [91](#)
 - job failures [146](#)
 - procedures for [157](#)
 - protection relationships [70](#)
 - Remote User option does not display [157](#)
 - storage failover link down condition [44, 46](#)
 - VMware Tools [157](#)
 - volume offline condition [47](#)
- types
 - of annotations [91](#)
 - of users [148](#)
- types of users
 - database users [15](#)

U

- Unified Manager servers
 - troubleshooting issue with installing or regenerating HTTPS certificates on when performance monitoring is enabled [159](#)
- Unified Manager web UI
 - overview [7](#)
- unmanaged relationships
 - editing lag thresholds settings for [37](#)

- user roles
 - assigning [37](#)
- users
 - adding [37](#)
 - capabilities associated with [148](#)
 - changing password for local [41](#)
 - creating [37](#)
 - creating a user with the Event Publisher role [53](#)
 - database [15](#)
 - diagnostic user [151](#)
 - maintenance user [150](#)
 - roles [147](#)
 - types [148](#)

V

- viewing
 - discovery status of clusters [40, 59](#)
 - notes about events [96](#)
- virtual appliances
 - overview of backup and restore process [18](#)
- virtual machine console
 - accessing the maintenance console [152](#)
- virtual machines
 - changing maintenance user password [154](#)
 - rebooting [154](#)
 - restarting [31](#)
 - shutting down [154](#)
- VM console
 - accessing the maintenance console [152](#)
- VMware Tools
 - troubleshooting [157](#)
- Volume details page
 - aborting SnapMirror relationships from [78](#)
 - creating SnapMirror relationships using [72](#)
 - creating SnapVault relationships using [73](#)
 - performing a reverse resync from [78](#)
 - removing relationships from [80](#)
 - restoring data from [87](#)
 - resynchronizing relationships from [80](#)
- volume offline
 - determining if caused by a down host cluster node [49](#)
 - determining if caused by a stopped SVM resulting from a down cluster node [50](#)
 - determining if caused by broken RAID disks [52](#)
 - troubleshooting [47](#)
- volumes
 - capacity information [102](#)
 - configuration information [102](#)

- configuring global threshold values for [36](#)
- creating rules to annotate [89](#)
- details about [102](#)
- efficiency information [102](#)
- FlexVol, defined [12](#)
- how they work [11](#)
- Infinite Volume defined [13](#)
- protection information [102](#)
- provisioning using SVM associations [15](#)
- SVMs with FlexVol, explained [9](#)
- SVMs with Infinite, explained [9](#)
- viewing annotations for [90](#)

Volumes page

- restoring data from [88](#)

Vservers

- See* SVMs

W

what the diagnostic user does

- defined [151](#)

Windows client

- using to retrieve the support bundle [93](#)

WinSCP

- using to retrieve support bundles [93](#)

Workflow Automation

- pairing with Unified Manager [71](#)

workflows

- incident analysis [55](#)
- performance monitoring [55](#)
- reference information common to [96](#)
- resolving lag issues [85](#)
- summary of common administrative [25](#)
- tasks common to [96](#)

workloads

- definition [21](#)
- system [21](#)
- user-defined [21](#)