

OnCommand Unified Manager Documentation

OnCommand Unified Manager 9.5

NetApp October 23, 2024

This PDF was generated from https://docs.netapp.com/us-en/oncommand-unified-manager-95/index.html on October 23, 2024. Always check docs.netapp.com for the latest.

Table of Contents

OnCommand Unified Manager Documentation	1
Release notes	2
Get started	3
Quick start instructions for VMWare installations	
Quick start instructions for Linux installations	3
Quick start instructions for Windows installations	5
Install Unified Manager	6
Introduction to OnCommand Unified Manager	6
Requirements for installing Unified Manager	7
Installing, upgrading, and removing Unified Manager software on VMware vSphere	17
Installing, upgrading, and removing Unified Manager software on Red Hat or CentOS	25
Installing, upgrading, and removing Unified Manager software on Windows	46
Perform configuration and administrative tasks	58
Configuring Unified Manager	58
Online Help	85
Introduction to OnCommand Unified Manager	85
Understanding the user interface	91
Monitoring cluster health and performance from the dashboards	99
Managing storage objects using the Favorites option	107
Managing events	
Managing alerts	195
Managing scripts	208
Managing health thresholds	
Managing user-defined performance thresholds	
Analyzing performance events	248
Resolving performance events	
Managing quotas	287
Managing and monitoring clusters and cluster object health	294
Managing and monitoring MetroCluster configurations	392
Managing annotations for storage objects	
Managing and monitoring groups	
Managing and monitoring protection relationships.	439
Executing protection workflows using OnCommand Workflow Automation	
Managing performance using performance capacity and available IOPS information	
Monitoring cluster performance from the Performance Cluster Landing page	528
Monitoring performance using the Performance Inventory pages	
Monitoring performance using the Performance Explorer pages	
Viewing object configuration information	
Understanding and using the Node Failover Planning page	599
Collecting data and monitoring workload performance	603
Analyzing workload performance	618
Managing reports	
Configuring backup and restore operations	682

Using Unified Manager REST APIs	690
Managing and monitoring Infinite Volumes	
Managing clusters	
Managing user access	
Managing authentication	
Managing security certificates	
Troubleshooting	
Monitor and manage cluster performance	
Introduction to OnCommand Unified Manager performance monitoring	
Navigating performance workflows in the Unified Manager GUI	
Understanding performance events and alerts	
Managing user-defined performance thresholds	783
Monitoring cluster performance from the Performance Dashboard	
Monitoring cluster performance from the Performance Cluster Landing page	
Monitoring performance using the Performance Inventory pages	
Monitoring performance using the Performance Explorer pages	807
Managing performance using performance capacity and available IOPS information	829
Understanding and using the Node Failover Planning page	837
Collecting data and monitoring workload performance	840
Analyzing workload performance	855
Analyzing performance events	864
Setting up a connection between a Unified Manager server and an external data provider	879
Monitor and manage cluster health	884
Introduction to OnCommand Unified Manager health monitoring	884
Common Unified Manager health workflows and tasks	886
Using the maintenance console	1057
Legal notices	1068
Copyright	1068
Trademarks	1068
Patents	1068
Privacy policy	1068
Open source	1068



Release notes

Provides a summary of new features, limitations, and known issues for OnCommand Unified Manager 9.5.

For more information, see the OnCommand Unified Manager Release Notes.

Get started

Quick start instructions for VMWare installations

System requirements

Operating system: VMware ESXi 5.5, 6.0, and 6.5

• RAM: 12 GB

CPU: 9572 MHz total

• Free disk space: 5 GB (thin provisioned), 152 GB (thick provisioned)

For detailed system requirements, see VMware software and installation requirements and Interoperability Matrix.

Installing OnCommand Unified Manager

Download the installer

- 1. Download the OnCommandUnifiedManager-9.5RC1.ovas installation package.
- 2. Save the file to a local directory or network directory that is accessible to your vSphere Client.

Install Unified Manager

VMware Tools is not included in the Unified Manager installation package. You must mount a CD-ROM or ISO image to install it as part of the Unified Manager installation process.

- 1. In the vSphere Client, click **File > Deploy OVF Template**.
- 2. Locate the OVA file and use the wizard to deploy the virtual appliance on the ESXi server.
- 3. On the Network Configuration page in the Properties tab, populate the fields as required for the type of installation you are performing:
 - For a static configuration; Enter required information in all fields (Secondary DNS is not required).
 - For DHCP using IPv4; Leave all the fields blank.
 - For DHCP using IPv6; Check the "Enable Auto IPv6 addressing" box and leave all the other fields blank.
- 4. Power on the VM.
- 5. Click the Console tab to view the initial boot process.
- 6. Follow the prompt to install VMware Tools on the VM.
- 7. Configure the time zone.
- 8. Enter a Unified Manager maintenance user name and password.

At the end of the installation, the information to connect to the Unified Manager web UI is displayed.

Quick start instructions for Linux installations

System requirements

 Operating system: Red Hat Enterprise Linux or CentOS 64-bit version 7.x architecture, installed using the "Server with GUI" base environment from the Software Selection option of the OS installer

• RAM: 12 GB

• CPU: 9572 MHz total

• Free disk space: 100 GB of disk space for /opt, 50 GB for the root partition

For detailed system requirements, see Red Hat Enterprise Linux and CentOS software and installation requirements and the Interoperability Matrix.

Installing OnCommand Unified Manager

Download the installer

- 1. Download the OnCommandUnifiedManager-rhel7-9.5RC1.zip installation package.
- 2. In the directory where you have downloaded the installation file, run:

```
# unzip OnCommandUnifiedManager-rhel7-9.5RC1.zip
```

Verify repository configuration

The procedures for configuring Red Hat Enterprise Linux or CentOS repositories are site specific. The pre_install_check.sh script included in the installation package can be optionally used to verify whether your operating system is correctly configured. If your system is connected to the internet, you automatically receive the instructions for setting up the Red Hat Enterprise Linux and MySQL repositories. For information on how to install on a system that has no internet connectivity, see Installing, upgrading, and removing Unified Manager software on Red Hat or CentOS.

```
# ./pre install check.sh
```

Install Unified Manager

Unified Manager uses the <code>yum</code> utility to install the software and any dependent software. As there are varying images of Red Hat Enterprise Linux or CentOS in different organizations, the packages installed depend on the software present in the images. The <code>yum</code> utility will determine the dependent software packages for installation. If you need more information on the dependent software packages, see <code>Installing</code>, upgrading, and removing <code>Unified Manager software on Red Hat or CentOS</code>.

As the root user, or using sudo, run the following command from the directory where the installation file was unzipped:

```
# yum install *.rpm
or
% sudo yum install *.rpms
```

At the end of the installation, the information to connect to the Unified Manager web UI is displayed. If you are unable to connect to the web UI, you may have to whitelist port 443. Contact your technical support team for

Quick start instructions for Windows installations

System requirements

- Operating Systems: Microsoft Windows Server 2012, 2012 R2, and 2016 64-bit Standard and Datacenter Editions
- RAM: 12 GB
- CPU: 9572 MHz total
- Free disk space: 100 GB of disk space for the installation directory, 50 GB of disk space for the MySQL data directory

For detailed system requirements, see Windows software and installation requirements and Interoperability Matrix.

Installing OnCommand Unified Manager

Download the installer

- 1. Download the OnCommandUnifiedManager-9.5RC1.exes installation package.
- 2. Copy the installation file to a directory on the target system.

Install Unified Manager

Microsoft .NET 4.5.2, or greater, must be installed. Unified Manager installs other required third-party packages as part of the installation. If you need more information on the dependent software packages, see Installing, upgrading, and removing Unified Manager software on Windows.

- 1. Log in to Windows using the default local administrator account.
- 2. In the directory where you downloaded the installation file, right-click and run the Unified Manager executable (.exe) file as an administrator.
- 3. When prompted, enter the user name and password to create the Unified Manager maintenance user.
- In the Database Connection wizard, enter the MySQL root password.
- 5. Follow the remaining prompts to complete the installation.
- 6. Click Finish at the end of the installation and the Unified Manager web UI is displayed.

Install Unified Manager

Introduction to OnCommand Unified Manager

OnCommand Unified Manager enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface. You can deploy Unified Manager on a Linux server, on a Windows server, or as a virtual appliance on a VMware host.

After you have completed the installation and have added the clusters that you want to manage, Unified Manager provides a graphical interface that displays the capacity, availability, protection, and performance status of the monitored storage systems.

Related information

NetApp Interoperability Matrix Tool

What the Unified Manager server does

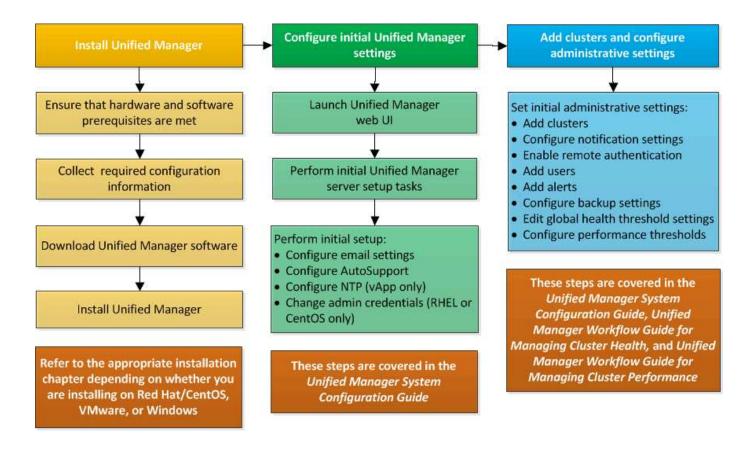
The Unified Manager server infrastructure consists of a data collection unit, a database, and an application server. It provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging.

Unified Manager collects cluster information, stores the data in the database, and analyzes the data to see if there are any cluster issues.

Overview of the installation sequence

The installation workflow describes the tasks that you must perform before you can use Unified Manager.

The chapters of this installation guide describe each of the items shown in the workflow below.



Requirements for installing Unified Manager

Before you can install Unified Manager you must ensure that the server on which you plan to install Unified Manager meets specific software, hardware, CPU, and memory requirements.

Related information

NetApp Interoperability Matrix Tool

Virtual infrastructure and hardware system requirements

Depending on whether you are installing Unified Manager on virtual infrastructure or on a physical system, it must meet minimum requirements for memory, CPU, and disk space.

The following table displays the values that are recommended for memory, CPU, and disk space resources. These values have been qualified so that Unified Manager meets acceptable performance levels.

Hardware configuration	Recommended settings
RAM	12 GB (minimum requirement 8 GB)
Processors	4 CPUs
CPU cycle capacity	9572 MHz total (minimum requirement 9572 MHz)

Hardward	e configuration	Recommended settings
Free disk	space	VMware:5 GB (thin provisioned)152 GB (thick provisioned)
	or CentOS: 150 GB, where the capacity is as follows:	Windows: 150 GB, where the capacity is allocated as follows:
• 50 GE	3 allotted to the root partition	100 GB of disk space for the installation directory
/opt	GB of free disk space allotted to the /netapp/data directory, which is mounted LVM drive or on a separate local disk ned to the target system	50 GB of disk space for the MySQL data directory
i	The /tmp directory should have at least 10 GB of free space and the /var/log directory should have at least 16 GB of free space.	

Unified Manager can be installed on systems with a small amount of memory, but the recommended 12 GB of RAM ensures that enough memory is available for optimal performance, and so that the system can accommodate additional clusters and storage objects as your configuration grows. You must not set any memory limits on the VM where Unified Manager is deployed, and you must not enable any features (for example, ballooning) that hinder the software from utilizing the allocated memory on the system.

Additionally, there is a limit to the number of nodes that a single instance of Unified Manager can monitor before you need to install a second instance of Unified Manager. See the *Best Practices Guide* for more details.

Technical Report 4621: Unified Manager Best Practices Guide

Memory-page swapping negatively impacts the performance of the system and the management application. Competing for CPU resources that are unavailable because of overall host utilization can degrade performance.

Dedicated use requirement

The physical or virtual system on which you install Unified Manager must be used exclusively for Unified Manager and must not be shared with other applications. Other applications might consume system resources and can drastically reduce the performance of Unified Manager.

Space requirements for backups

If you plan to use the Unified Manager backup and restore feature, you must allocate additional capacity so that the "data" directory or disk has 150 GB of space. A backup can be written to a local destination or to a remote destination. The best practice is to identify a remote location that is external to the Unified Manager host system that has a minimum of 150 GB of space.

Host connectivity requirements

The physical system or virtual system on which you install Unified Manager must be configured in such a way that you can successfully ping the host name from the host itself. In case of IPv6 configuration, you should verify that ping6 to the host name is successful to ensure that the Unified Manager installation succeeds.

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 designated a name for the network host. If you configured the network using DHCP, you should obtain the host name from the DNS.

If you plan to allow users to access Unified Manager by using the short name instead of using the fully qualified domain name (FQDN) or IP address, then your network configuration has to resolve this short name to a valid FQDN.

Mounted /opt/netapp or /opt/netapp/data requirements

You can mount /opt/netapp or /opt/netapp/data on an NAS or SAN device. Note that using remote mount points may cause scaling issues. If you do use a remote mount point, ensure that your SAN or NAS network has sufficient capacity to meet the I/O needs of Unified Manager. This capacity will vary and may increase based on the number of clusters and storage objects you are monitoring.

If you have mounted <code>/opt/netapp</code> or <code>/opt/netapp/data</code> from anywhere other that the root file system, and you have SELinux enabled in your environment, you must set the correct context for the mounted directories.

See the topic SELinux requirements for mounting /opt/netapp or /opt/netapp/data on an NFS or CIFS share for information about setting the correct SELinux context.

VMware software and installation requirements

The VMware vSphere system on which you install Unified Manager requires specific versions of the operating system and supporting software.

Operating system software

The following versions of VMware ESXi are supported:

• ESXi 5.5, 6.0, and 6.5

The following versions of vSphere are supported:

VMware vCenter Server 5.5, 6.0, and 6.5

See the Interoperability Matrix for the complete and most current list of supported ESXi versions.

mysupport.netapp.com/matrix

The VMware ESXi server time must be the same as the NTP server time for the virtual appliance to function correctly. Synchronizing the VMware ESXi server time with the NTP server time prevents a time failure.

Installation requirements

VMware High Availability for the Unified Manager virtual appliance is supported.

If you deploy an NFS datastore on a storage system that is running ONTAP software, you must use the NetApp NFS Plug-in for VMware VAAI to use thick provisioning.

If deployment fails using your High Availability-enabled environment because of insufficient resources, you may need to modify the Cluster Features Virtual Machine Options by disabling the VM Restart Priority, and leaving the Host Isolation Response powered on.

Red Hat Enterprise Linux and CentOS software and installation requirements

The Linux system on which you install Unified Manager requires specific versions of the operating system and supporting software.

Operating system software

The Linux system must have the following versions of the operating system and supporting software installed:

• Red Hat Enterprise Linux or CentOS 64-bit version 7.x

Red Hat Enterprise Linux 6.x is not supported starting with Unified Manager 9.4.

See the Interoperability Matrix for the complete and most current list of supported Red Hat Enterprise Linux and CentOS versions.

mysupport.netapp.com/matrix

The following third-party packages are required:

- MySQL Community Edition version 5.7.23 or later versions in the 5.7 family (from the MySQL repository)
- OpenJDK version 11 (from the Red Hat Extra Enterprise Linux Server repository)



Oracle Java is not supported starting with Unified Manager 9.5.

• p7zip version 16.02 or later (from the Red Hat Extra Packages for Enterprise Linux repository)



If you plan to upgrade any of the third-party software after Unified Manager has been running, you must shut down Unified Manager first. After the third-party software installation is complete, you can restart Unified Manager.

User authorization requirements

Installation of Unified Manager on a Red Hat Enterprise Linux system or CentOS system can be performed by the root user or by non-root users by using the sudo command.

Installation requirements

The best practices for installing Red Hat Enterprise Linux or CentOS and the associated repositories on your system are as follows:

- You must install Red Hat Enterprise Linux or CentOS according to Red Hat best practices, and you should select the following default options, which requires selecting "Server with GUI".
- While installing Unified Manager on Red Hat Enterprise Linux or CentOS, the system must have access to the appropriate repository so that the installation program can access and install all the required software

dependencies.

• For the yum installer to find dependent software in the Red Hat Enterprise Linux repositories, you must have registered the system during the Red Hat Enterprise Linux installation or afterwards by using a valid Red Hat subscription.

See the Red Hat documentation for information about the Red Hat Subscription Manager.

• You must enable the Extra Packages for Enterprise Linux (EPEL) repository to successfully install the required third-party utilities on your system.

If the EPEL repository is not configured on your system, you must manually download and configure the repository.

Manually configuring the EPEL repository

• If the correct version of MySQL is not installed, you must enable the MySQL repository to successfully install MySQL software on your system.

If the MySQL repository is not configured on your system, you must manually download and configure the repository.

Manually configuring the MySQL repository

If your system does not have internet access, and the repositories are not mirrored from an internet-connected system to the unconnected system, you should follow the installation instructions to determine the external software dependencies of your system. Then you can download the required software to the internet-connected system, and copy the <code>.rpm</code> files to the system on which you plan to install Unified Manager. To download the artifacts and packages, you must use the <code>yum install</code> command. You must ensure that the two systems are running the same operating system version and that the subscription license is for the appropriate Red Hat Enterprise Linux or CentOS version.



You must not install the required third-party software from repositories other than the repositories that are listed here. Software installed from the Red Hat repositories is designed explicitly for Red Hat Enterprise Linux, and conforms to Red Hat best practices (directory layouts, permissions, and so on). Software from other locations might not follow these guidelines, which might cause the Unified Manager installation to fail, or might cause issues with future upgrades.

Port 443 requirement

Generic images from Red Hat and CentOS block external access to port 443. If your browser is unable to connect to your OnCommand product, this may be the issue. The following command enables access to port 443 for all external users and applications: # firewall-cmd -zone=public -add-port=443/tcp -permanent; firewall-cmd -reload

Consult with your IT department prior to executing this command to see if your security policies require a different procedure.

Windows software and installation requirements

For the successful installation of Unified Manager on Windows, you must ensure that the system on which Unified Manager is being installed meets the software requirements.

Operating system software

Unified Manager runs only on a 64-bit English language Windows operating system. You can install Unified Manager on the following Windows platforms:

- Microsoft Windows Server 2012 Standard and Datacenter Edition
- Microsoft Windows Server 2012 R2 Standard and Datacenter Edition
- Microsoft Windows Server 2016 Standard and Datacenter Edition.



On Windows Server 2012 R2, Windows update KB2919355 must be installed on the target system or the installation will fail.

Note that Windows Server 2008 is not supported as it was in earlier releases. See the Interoperability Matrix for the complete and most current list of supported Windows versions.

mysupport.netapp.com/matrix

The server should be dedicated to running Unified Manager; no other applications should be installed on the server.

The following third-party packages are required:

- Microsoft Visual C++ 2015 Redistributable package version 14.0.24212
- Microsoft Visual C++ Redistributable Packages for Visual Studio 2013 version 12.0.40660
- MySQL Community Edition version 5.7.23, or later versions in the 5.7 family
- OpenJDK version 11
- p7zip version 18.01 or later

If these third-party packages are not installed, Unified Manager installs them as part of the installation.



Starting with Unified Manager 9.5, OpenJDK is provided in the Unified Manager installation package and installed automatically. Oracle Java is not supported starting with Unified Manager 9.5.

If MySQL is pre-installed, you must ensure that:

- It is using the default port.
- The sample databases are not installed.
- The service name is "MYSQL".



If you plan to upgrade any of the third-party software after Unified Manager has been running, you must shut down Unified Manager first. After the third-party software installation is complete you can restart Unified Manager.

Installation requirements

- Microsoft .NET 4.5.2, or greater, must be installed.
- You must reserve 2 GB of disk space for the temp directory to extract the installation files.
- You must reserve 2 GB of disk space in the Windows drive for caching the Unified Manager MSI files.

- The Microsoft Windows Server on which you want to install Unified Manager must be configured with a fully qualified domain name (FQDN) such that ping responses to the host name and FQDN are successful.
- You must disable Microsoft IIS worldwide web publishing service and ensure that ports 80 and 443 are free.
- You must make sure that the Remote Desktop Session Host setting for "Windows Installer RDS Compatibility" is disabled during the installation.
- UDP port 514 must be free, and must not be used by any other service.

The Unified Manager installation program configures the following exclusions in Windows Defender:



- Unified Manager data directory (Windows Server 2016 only)
- · Unified Manager installation directory
- · MySQL data directory

If your server has a different antivirus scanner installed you must configure these exclusions manually.

Supported browsers

To access the Unified Manager UI, you must use a supported browser.

Unified Manager has been tested with the following browsers; other browsers might work but have not been qualified. See the Interoperability Matrix for the complete list of supported browser versions.

mysupport.netapp.com/matrix

- Mozilla Firefox ESR 60
- · Google Chrome version 68 and 69
- Microsoft Internet Explorer 11

For all browsers, disabling popup blockers helps ensure that software features display properly.

For Internet Explorer, you must ensure that Compatibility View is disabled, and Document Mode is set to the default. See the Microsoft IE documentation for information about these settings.



Firefox and Chrome are the preferred browsers as there have been some cases where complex UI pages load more slowly when using Internet Explorer.

If you are planning to configure Unified Manager for SAML authentication so that an identity provider (IdP) authenticates users, check the list of browsers supported by the IdP as well.

Protocol and port requirements

Using a browser, API client, or SSH, the required ports must be accessible to the Unified Manager UI and APIs. The required ports and protocols enable communication between the Unified Manager server and the managed storage systems, servers, and other components.

Connections to the Unified Manager server

You do not have to specify port numbers when connecting to the Unified Manager web UI, because default ports are always used. For example, because Unified Manager always runs on its default port, you can enter https://<host>instead of https://<host>: 443. The default port numbers cannot be changed.

The Unified Manager server uses specific protocols to access the following interfaces:

Interface	Protocol	Port	Description
Unified Manager web UI	HTTP	80	Used to access the Unified Manager web UI; automatically redirects to the secure port 443.
Unified Manager web UI and programs using APIs	HTTPS	443	Used to securely access the Unified Manager web UI or to make API calls; API calls can only be made using HTTPS.
Maintenance console	SSH/SFTP	22	Used to access the maintenance console and retrieve support bundles.
Linux command line	SSH/SFTP	22	Used to access the Red Hat Enterprise Linux or CentOS command line and retrieve support bundles.
MySQL database	MySQL	3306	Used to enable OnCommand Workflow Automation and OnCommand API Services access to Unified Manager.
Syslog	UDP	514	Used to access subscription-based EMS messages from ONTAP systems and to create events based on the messages.
REST	HTTPS	9443	Used to access realtime REST API-based EMS events from authenticated ONTAP systems.

Connections from the Unified Manager server

You must configure your firewall to open ports that enable communication between the Unified Manager server and managed storage systems, servers, and other components. If a port is not open, communication fails.

Depending on your environment, you can choose to modify the ports and protocols used by the Unified Manager server to connect to specific destinations.

The Unified Manager server connects using the following protocols and ports to the managed storage systems, servers, and other components:

Destination	Protocol	Port	Description
Storage system	HTTPS	443/TCP	Used to monitor and manage storage systems.
Storage system	NDMP	10000/TCP	Used for certain Snapshot restore operations.
AutoSupport server	HTTPS	443	Used to send AutoSupport information. Requires Internet access to perform this function.
Authentication server	LDAP	389	Used to make authentication requests, and user and group lookup requests.
LDAPS	636	Used for secure LDAP communication.	Mail server
SMTP	25	Used to send alert notification emails.	SNMP trap sender
SNMPv1 or SNMPv3	162/UDP	Used to send alert notification SNMP traps.	External data provider server
TCP	2003	Used to send performance data to an external data provider, such as Graphite.	NTP server

Completing the worksheet

Before you install and configure Unified Manager, you should have specific information about your environment readily available. You can record the information in the worksheet.

Unified Manager installation information

The details required to install Unified Manager.

System on which software is deployed	Your value
ESXi server IP address (VMware only)	
Host fully qualified domain name	
Host IP address	
Network mask	
Gateway IP address	
Primary DNS address	
Secondary DNS address	
Search domains	
Maintenance user name	
Maintenance user password	

Unified Manager configuration information

The details to configure Unified Manager after installation. Some values are optional depending on your configuration.

Setting	Your value
Maintenance user email address	
NTP server (VMware only)	
SMTP server host name or IP address	
SMTP user name	
SMTP password	
SMTP port	25 (Default value)
Email from which alert notifications are sent	

Setting	Your value
Authentication server host name or IP address	
Active Directory administrator name or LDAP bind distinguished name	
Active Directory password or LDAP bind password	
Authentication server base distinguished name	
Identity provider (IdP) URL	
Identity provider (IdP) metadata	
SNMP trap destination host IP address	
SNMP port	

Cluster information

The details for the storage systems that you will manage using Unified Manager.

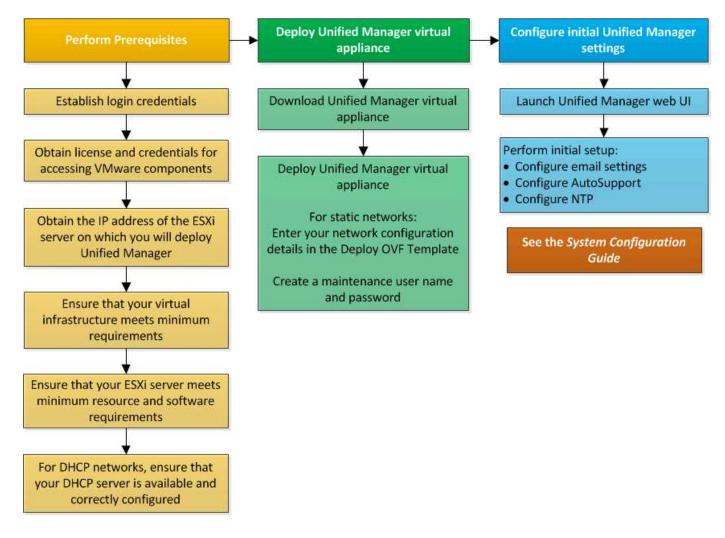
Cluster 1	of N	Your value
Host nam	e or cluster-management IP address	
ONTAP a	dministrator user name	
i	The administrator must have been assigned the "admin" role.	
ONTAP a	dministrator password	
Protocol ((HTTP or HTTPS)	

Installing, upgrading, and removing Unified Manager software on VMware vSphere

On VMware vSphere systems, you can install Unified Manager software, upgrade to a newer version of software, or remove the Unified Manager virtual appliance.

Overview of the deployment process on VMware

The deployment workflow describes the tasks that you must perform before you can use Unified Manager.



Deploying Unified Manager

Deploying Unified Manager includes downloading software, deploying the virtual appliance, creating a maintenance user name and password, and performing the initial setup in the web UI.

Before you begin

· You must have completed the system requirements for deployment.

System requirements

- · You must have the following information:
 - Login credentials for the NetApp Support Site
 - Credentials for accessing the VMware vCenter Server and vSphere Web Client (for vSphere version 6.5) or vSphere Client (for vSphere version 5.5 or 6.0)
 - IP address of the ESXi server on which you are deploying the Unified Manager virtual appliance
 - · Details about the data center, such as storage space in the datastore and memory requirements
 - IPv6 must be enabled on the host if you are planning to use IPv6 addressing.
 - CD-ROM or ISO image of VMware Tools

About this task

You can deploy Unified Manager as a virtual appliance on a VMware ESXi server.

You must access the maintenance console by using the VMware console, and not by using SSH.

VMware Tools are not included in the Unified Manager.ova file, and must be installed separately.

After you finish

After finishing the deployment and initial setup, you can either add clusters, or configure additional network settings in the maintenance console, and then access the web UI.

Downloading the Unified Manager OVA file

You must download the Unified Manager.ova file from the NetApp Support Site to deploy Unified Manager as a virtual appliance.

Before you begin

You must have login credentials for the NetApp Support Site.

About this task

The .ova file contains the Unified Manager software configured in a virtual appliance.

Steps

1. Log in to the NetApp Support Site, and navigate to the Download page for installing Unified Manager on the VMware vSphere.

NetApp Downloads: Software

- 2. Download and save the .ova file to a local directory or network directory that is accessible to your vSphere Client.
- 3. Verify the checksum to ensure that the software downloaded correctly.

Deploying the Unified Manager virtual appliance

You can deploy the Unified Manager virtual appliance after you download the .ova file from the NetApp Support Site. You must use the vSphere Web Client (for vSphere version 6.5) or vSphere Client (for vSphere version 5.5 or 6.0) to deploy the virtual appliance on an ESXi server. When you deploy the virtual appliance, a virtual machine is created.

Before you begin

You must have reviewed the system requirements. If changes are required to meet the system requirements, you must implement the changes before deploying the Unified Manager virtual appliance.

Virtual infrastructure requirements

VMware software and installation requirements

If you use DHCP, you must ensure that the DHCP server is available, and that the DHCP and virtual machine (VM) network adapter configurations are correct. DHCP is configured by default.

If you use a static networking configuration, you must ensure that the IP address is not duplicated in the same subnet, and that the appropriate DNS server entries have been configured.

You must have the following information before deploying the virtual appliance:

- Credentials for accessing the VMware vCenter Server and vSphere Web Client (for vSphere version 6.5) or vSphere Client (for vSphere version 5.5 or 6.0)
- IP address of the ESXi server on which you are deploying the Unified Manager virtual appliance
- · Details about the data center, such as availability of storage space
- If you are not using DHCP, you must have the IPv4 or IPv6 addresses for the networking devices to which you are planning to connect:
 - Fully qualified domain name (FQDN) of the host
 - IP address of the host
 - Network mask
 - IP address of the default gateway
 - Primary and secondary DNS addresses
 - Search domains
- CD-ROM or ISO image for the VMware Tools

About this task

VMware Tools are not included in the .ova file. You must install the VMware Tools separately.

When the virtual appliance is deployed, a unique self-signed certificate for HTTPS access is generated. When accessing the Unified Manager web UI, you might see a browser warning about untrusted certificates.

VMware High Availability for the Unified Manager virtual appliance is supported.

Steps

- 1. In the vSphere Client, click File > Deploy OVF Template.
- 2. Complete the **Deploy OVF Template** wizard to deploy the Unified Manager virtual appliance.

On the Networking Configuration page:

- Leave all the fields blank when using DHCP and IPv4 addressing.
- Check the "Enable Auto IPv6 addressing" box, and leave all the other fields blank when using DHCP and IPv6 addressing.
- If you want to use a static network configuration, you can complete the fields on this page and these
 settings are applied during deployment. You must ensure that the IP address is unique to the host on
 which it is deployed, that it is not already in use, and that it has a valid DNS entry.
- 3. After the Unified Manager virtual appliance is deployed to the ESXi server, power on the VM by right-clicking the VM, and then selecting **Power On**.

If the Power On operation fails because of insufficient resources, you must add resources and then retry

the installation.

4. Click the Console tab.

The initial boot process takes a few minutes to complete.

5. Follow the prompt to install the VMware Tools on the VM.

When using the vSphere Web Client with vSphere 6.5 you need to manually mount the VMware Tools ISO image. From the VM you need to select **Edit Settings > Virtual Hardware > CD/DVD drive x > Datastore ISO file** and then click **Browse** to select the file linux.iso as the mount image.

To configure your time zone, enter your geographic area and your city or region as prompted in the VM Console window.

All the date information that is displayed uses the time zone that is configured for Unified Manager, regardless of the time zone setting on your managed devices. You should be aware of this when comparing time stamps. If your storage systems and the management server are configured with the same NTP server, they refer to the same instant in time, even if they appear differently. For example, if you create a Snapshot copy using a device that is configured using a different time zone than that of the management server, the time reflected in the time stamp is the management server time.

7. If no DHCP services are available, or if there is an error in the details for the static network configuration, select one of the following options:

If you use	Then do this
DHCP	Select Retry DHCP . If you plan to use DHCP, you should ensure that it is configured correctly. If you use a DHCP-enabled network, the FQDN and DNS server entries are given to the virtual appliance automatically. If DHCP is not properly configured with DNS, the host name "OnCommand" is automatically assigned and associated with the security certificate. If you have not set up a DHCP-enabled network, you must manually enter the networking configuration information.
A static network configuration	 a. Select Enter the details for static network configuration. The configuration process takes a few minutes to complete. b. Confirm the values that you entered, and select Y.

8. At the prompt, enter a maintenance user name, and click **Enter**.

The maintenance user name must start with a letter from a-z, followed by any combination of -, a-z, or 0-9.

9. At the prompt, enter a password, and click **Enter**.

The VM console displays the URL for the Unified Manager web UI.

After you finish

You can access the web UI to perform the initial setup of Unified Manager, as described in the *OnCommand Unified Manager System Configuration Guide*.

Upgrading Unified Manager on VMware

You can upgrade to Unified Manager version 9.5 only from instances of Unified Manager 7.3 or 9.4.

About this task

During the upgrade process, Unified Manager is unavailable. You should complete any running operations before upgrading Unified Manager.

If Unified Manager is paired with an instance of OnCommand Workflow Automation, and there are new versions of software available for both products, you must disconnect the two products and then set up a new Workflow Automation connection after performing the upgrades. If you are performing an upgrade to only one of the products, then you should log into Workflow Automation after the upgrade and verify that it is still acquiring data from Unified Manager.

Downloading the Unified Manager ISO image

Before upgrading Unified Manager, you must download the Unified Manager ISO image from the NetApp Support Site.

Before you begin

You must have login credentials for the NetApp Support Site.

Steps

- 1. Log in to the NetApp Support Site and navigate to the Software Download page.
- 2. Download and save the .iso image file to a local directory or network directory that is accessible to your vSphere Client.
- 3. Verify the checksum to ensure that the software downloaded correctly.

Related information

NetApp Support

Upgrading the Unified Manager virtual appliance

You can upgrade from Unified Manager version 7.3 or 9.4 to Unified Manager 9.5.

Before you begin

- You must have downloaded the .iso file from the NetApp Support Site.
- The system on which you are upgrading Unified Manager must meet the system and software requirements.

Virtual infrastructure requirements

VMware software and installation requirements

- For vSphere 6.5 users, you must have installed the VMware Remote Console (VMRC).
- · You must have the following information:
 - Login credentials for the NetApp Support Site
 - Credentials for accessing the VMware vCenter Server and vSphere Web Client (for vSphere version 6.5) or vSphere Client (for vSphere version 5.5 or 6.0)
 - · Credentials for the Unified Manager maintenance user

About this task

During the upgrade process, Unified Manager is unavailable. You should complete any running operations before upgrading Unified Manager.

If you have paired Workflow Automation and Unified Manager, you must manually update the host name in Workflow Automation.

Steps

- 1. In the vSphere Client, click **Home > Inventory > VMs and Templates**.
- Select the virtual machine (VM) on which the Unified Manager virtual appliance is installed.
- 3. If the Unified Manager VM is running, navigate to Summary > Commands > Shut Down Guest.
- 4. Create a backup copy—such as a snapshot or clone—of the Unified Manager VM to create an application-consistent backup.
- 5. From the vSphere Client, power on the Unified Manager VM.
- 6. Select the Unified Manager upgrade image:

If you are using	Then do this
vSphere 5.5 or 6.0	a. Click the CD/DVD Drive icon, and select Connect to ISO image on local disk.
	b. Select the OnCommandUnifiedManager- 9.5-virtual-update.iso file, and click Open.
vSphere 6.5	a. Launch the VMware Remote Console.
	 b. Click the CDROM icon, and select Connect to Disk Image File (.iso).
	c. Select the OnCommandUnifiedManager- 9.5-virtual-update.iso file, and click Open.

- 7. Click the Console tab.
- 8. Log in to the Unified Manager maintenance console.
- 9. In the Main Menu, select Upgrade.

A message is displayed that Unified Manager will be unavailable during the upgrade process, and will resume after completion.

10. Type y to continue.

A warning is displayed, reminding you to back up the virtual machine on which the virtual appliance resides.

11. Type y to continue.

The upgrade process and the restart of Unified Manager services can take several minutes to complete.

12. Press any key to continue.

You are automatically logged out of the maintenance console.

13. Log in to the maintenance console, and verify the version of Unified Manager.

After you finish

You can log in to the web UI to use the upgraded version of Unified Manager. Note that you must wait for the discovery process to finish before performing any task in the UI.

Restarting the Unified Manager virtual machine

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

Before you begin

- The virtual appliance must be powered on.
- You must be logged in to the Unified Manager maintenance console as the maintenance user.

About this task

You can also restart the virtual machine from vSphere by using the VMware Restart Guest option.

Steps

- 1. In the maintenance console, select **System Configuration > Reboot Virtual Machine**.
- 2. Start the Unified Manager graphical user interface (GUI) from your browser, and log in.

Related information

VMware vSphere PowerCLI Cmdlets Reference: Restart-VMGuest

Removing Unified Manager from VMware

You can uninstall Unified Manager by destroying the virtual appliance on which the Unified Manager software is installed.

Before you begin

- You must have credentials for accessing VMware vCenter Server and vSphere Web Client (for vSphere version 6.5) or vSphere Client (for vSphere version 5.5 or 6.0).
- The Unified Manager server must not have an active connection to an external data provider.

If there is an active connection, you must delete the connection by using the Unified Managermaintenance console.

- The Unified Manager server must not have an active connection to a Workflow Automation server.
 - If there is an active connection, you must delete the connection by using the Administration menu.
- All clusters (data sources) must be removed from the Unified Manager server before you delete the virtual machine (VM).

Steps

- 1. Use the Unified Managermaintenance console to verify that the Unified Manager server does not have an active connection to an external data provider.
- 2. In the vSphere Client, click Home > Inventory > VMs and Templates.
- 3. Select the VM that you want to destroy, and click the **Summary** tab.
- If the VM is running, click Power > Shut Down Guest.
- 5. Right-click the VM that you want to destroy, and click Delete from Disk.

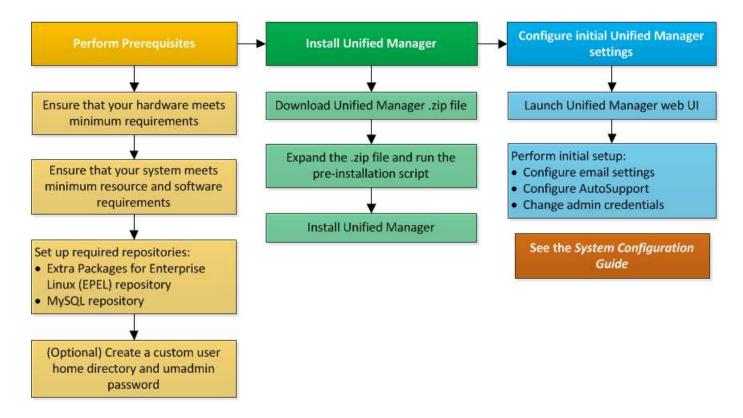
Installing, upgrading, and removing Unified Manager software on Red Hat or CentOS

On Linux systems, you can install Unified Manager software, upgrade to a newer version of software, or remove Unified Manager.

Unified Manager can be installed on Red Hat Enterprise Linux or CentOS servers. The Linux server on which you install Unified Manager can be running either on a physical machine or on a virtual machine running on VMware ESXi, Microsoft Hyper-V, or Citrix XenServer.

Overview of the installation process on Red Hat or CentOS

The installation workflow describes the tasks that you must perform before you can use Unified Manager.



Setting up required software repositories

The system must have access to certain repositories so that the installation program can access and install all required software dependencies.

Manually configuring the EPEL repository

If the system on which you are installing Unified Manager does not have access to the Extra Packages for Enterprise Linux (EPEL) repository, then you must manually download and configure the repository for a successful installation.

About this task

The EPEL repository provides access to the required third-party utilities that must be installed on your system. You use the EPEL repository whether you are installing Unified Manager on a Red Hat or CentOS system.

Steps

- 1. Download the EPEL repository for your installation: wget https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
- 2. Configure the EPEL repository: yum install epel-release-latest-7.noarch.rpm

Manually configuring the MySQL repository

If the system on which you are installing Unified Manager does not have access to the MySQL Community Edition repository, then you must manually download and configure the repository for a successful installation.

About this task

The MySQL repository provides access to the required MySQL software that must be installed on your system.



This task will fail if the system does not have Internet connectivity. Refer to the MySQL documentation if the system on which you are installing Unified Manager does not have Internet access.

Steps

- 1. Download the appropriate MySQL repository for your installation: wget http://repo.mysql.com/yum/mysql-5.7-community/el/7/x86_64/mysql57-community-release-el7-7.noarch.rpm
- 2. Configure the MySQL repository: yum install mysql57-community-release-el7-7.noarch.rpm

SELinux requirements for mounting /opt/netapp **or** /opt/netapp/data **on an NFS or CIFS share**

If you are planning to mount /opt/netapp or /opt/netapp/data on an NAS or SAN device, and you have SELinux enabled, you need to be aware of the following considerations.

About this task

If are planning to mount /opt/netapp or /opt/netapp/data from anywhere other that the root file system, and you have SELinux enabled in your environment, you must set the correct context for the mounted directories. Follow these two steps for setting and confirming the correct SELinux context.

- Configure SELinux context when /opt/netapp/data is mounted
- Configure SELinux context when /opt/netapp is mounted

Configuring the SELinux context when /opt/netapp/data is mounted

If you have mounted $\protect\operatorname{\mathsf{Mopt/netapp/data}}$ in your system and SELinux is set to $\protect\operatorname{\mathsf{Enforcing}}$, ensure that the SELinux context type for $\protect\operatorname{\mathsf{Mopt/netapp/data}}$ is set to $\protect\operatorname{\mathsf{mysqld_db_t}}$, which is the default context element for the location of the database files.

1. Run this command to check the context: ls -dZ /opt/netapp/data

A sample output:

```
drwxr-xr-x. mysql root unconfined_u:object_r:default_t:s0
/opt/netapp/data
```

In this output, the context is default t that must be changed to mysqld db t.

- 2. Perform these steps to set the context, based on how you have mounted /opt/netapp/data.
 - a. Run the following commands to set the context to mysqld_db_t: semanage fcontext -a -t

```
mysql_db_t "/opt/netapp/data"``restorecon -R -v /opt/netapp/data
```

- b. If you have configured /opt/netapp/data in /etc/fstab, you must edit the /etc/fstab file. For the /opt/netapp/data/ mount option, add the MySQL label as: context=system u:object r:mysqld db t:s0
- c. Unmount and remount /opt/netapp/data/ for enabling the context.
- 3. Verify whether the context is set correctly: ls -dZ /opt/netapp/data/

```
drwxr-xr-x. mysql root unconfined_u:object_r:mysqld_db_t:s0
/opt/netapp/data/
```

Configuring the SELinux context when /opt/netapp is mounted

After setting the correct context for /opt/netapp/data/, ensure that the parent directory /opt/netapp does not have the SELinux context set to file t.

1. Run this command to check the context: ls -dZ /opt/netapp

A sample output:

```
drwxr-xr-x. mysql root unconfined_u:object_r:file_t:s0 /opt/netapp
```

In this output, the context is file_t that must be changed. The following commands set the context to usr_t. You can set the context to any value other than file_t based on your security requirements.

- 2. Perform these steps to set the context, based on how you have mounted /opt/netapp.
 - a. Run the following commands to set the context: semanage fcontext -a -t usr_t "/opt/netapp" ``restorecon -v /opt/netapp
 - b. If you have configured /opt/netapp in /etc/fstab, you must edit the /etc/fstab file. For the /opt/netapp mount option, add the MySQL label as: context=system_u:object_r:usr_t:s0
 - c. Unmount and remount /opt/netapp for enabling the context.
- 3. Verify whether the context is set correctly: ls -dZ /opt/netapp

```
drwxr-xr-x. mysql root unconfined_u:object_r:usr_t:s0 /opt/netapp
```

Installing Unified Manager on Red Hat Enterprise Linux or CentOS

It is important that you understand that the sequence of steps to download and install Unified Manager varies according to your installation scenario. Before you install Unified Manager on Red Hat Enterprise Linux or CentOS, you can decide if you want to configure Unified Manager for high availability.

Creating a custom user home directory and umadmin password prior to installation

You can create a custom home directory and define your own umadmin user password prior to installing Unified Manager. This task is optional, but some sites might need the flexibility to override Unified Manager installation default settings.

Before you begin

- The system must meet the requirements described in Hardware system requirements.
- You must be able to log in as the root user to the Red Hat Enterprise Linux or CentOS system.

About this task

The default Unified Manager installation performs the following tasks:

- Creates the umadmin user with /home/umadmin as the home directory.
- · Assigns the default password "admin" to the umadmin user.

Because some installation environments restrict access to /home, the installation fails. You must create the home directory in a different location. Additionally, some sites might have rules about password complexity or require that passwords be set by local administrators rather than being set by the installing program.

If your installation environment requires that you override these installation default settings, follow these steps to create a custom home directory and to define the umadmin user's password.

When this information is defined prior to installation, the installation script discovers these settings and uses the defined values instead of using the installation default settings.

Additionally, the default Unified Manager installation includes the umadmin user in the sudoers files (ocum_sudoers and ocie_sudoers) in the /etc/sudoers.d/ directory. If you remove this content from your environment because of security policies, or because of some security monitoring tool, you must add it back. You need to preserve the sudoers configuration because some Unified Manager operations require these sudo privileges.

Steps

- 1. Log in as the root user to the server.
- Create the umadmin group account called "maintenance": groupadd maintenance
- 3. Create the user account "umadmin" in the maintenance group under a home directory of your choice:adduser --home <home directory > -g maintenance umadmin
- 4. Define the umadmin password:passwd umadmin

The system prompts you to enter a new password string for the umadmin user.

After you finish

After you have installed Unified Manager you must specify the umadmin user login shell.

Downloading Unified Manager for Red Hat Enterprise Linux or CentOS

You must download the Unified Manager.zip file from the NetApp Support Site to install Unified Manager.

Before you begin

You must have login credentials for the NetApp Support Site.

About this task

You download the same Unified Manager installation package for both Red Hat Enterprise Linux and CentOS systems.

Steps

1. Log in to the NetApp Support Site, and navigate to the Download page for installing Unified Manager on the Red Hat Enterprise Linux platform.

NetApp Downloads: Software

- 2. Download the Unified Manager. zip file to a directory on the target system.
- 3. Verify the checksum to ensure that the software downloaded correctly.

Installing Unified Manager on Red Hat Enterprise Linux or CentOS

You can install Unified Manager on a physical or virtual Red Hat Enterprise Linux or CentOS platform.

Before you begin

 The system on which you want to install Unified Manager must meet the system and software requirements.

Hardware system requirements

Red Hat and CentOS software and installation requirements

- You must have downloaded the Unified Manager.zip file from the NetApp Support Site to the target system.
- You must have a supported web browser.
- · Your terminal emulation software must have scrollback enabled.

About this task

The Red Hat Enterprise Linux or CentOS system may have all the required versions of the required supporting software (Java, MySQL, additional utilities) installed, or it may have only some of the required software installed, or it may be a newly installed system with none of the required software installed.

Steps

- 1. Log in to the server on which you are installing Unified Manager.
- Enter the appropriate commands to assess what software might require installation or upgrade on the target system to support installation:

Required software and minimum version	Command to verify software and version	
OpenJDK version 11	java -version	
MySQL 5.7.23 Community Edition	rpm -qa grep -i mysql	
p7zip 9.20.1	rpm -qa grep p7zip	

3. If any version of the listed software is earlier than the required version, enter the appropriate command to uninstall that module:

Software to uninstall		Command to uninstall the software	
MySQL		rpm -e <mysql_package_name></mysql_package_name>	
i	Uninstall any version that is not MySQL 5.7.23 Community Edition or later.	i	If you receive dependency errors, you must add thenodeps option to uninstall the component.
All other modules		yum remove module_name	

4. Navigate to the directory where you downloaded the installation .zip file and expand the Unified Manager bundle: unzip OnCommandUnifiedManager-rhel7-9.5.zip

The required .rpm modules for Unified Manager are unzipped to the target directory.

5. Verify that the following modules are available in the directory: ls *.rpm

```
ocie-au-<version>.x86_64.rpm
ocie-server-<version>.x86_64.rpm
ocie-serverbase-<version>.x86_64.rpm
netapp-application-server-<version>.x86_64.rpm
netapp-platform-base-<version>.x86_64.rpm
netapp-ocum-<version>.x86_64.rpm
```

6. Run the pre-installation script to ensure that there are no system configuration settings or any installed software that will conflict with the installation of Unified Manager: pre install check.sh

The pre-installation script checks that the system has a valid Red Hat subscription, and that it has access to the required software repositories. If the script identifies any issues, you must fix the issues prior to installing Unified Manager.



You must perform step 7 *only* if you are required to manually download the packages that are required for your installation. If your system has Internet access and all the required packages are available, go to step 8.

- 7. For systems that are not connected to the Internet or that are not using the Red Hat Enterprise Linux repositories, perform the following steps to determine whether you are missing any required packages, and then download those packages:
 - a. On the system on which you are installing Unified Manager, view the list of available and unavailable packages: yum install *.rpm --assumeno

The items in the "Installing:" section are the packages that are available in the current directory, and the items in the "Installing for dependencies:" section are the packages that are missing on your system.

b. On a system that has Internet access, download the missing packages: yum install
 <package_name\> --downloadonly --downloaddir=.



Because the plug-in "yum-plugin-downloadonly" is not always enabled on Red Hat Enterprise Linux systems, you might need to enable the functionality to download a package without installing it: yum install yum-plugin-downloadonly

- c. Copy the missing packages from the Internet-connected system to your installation system.
- 8. Install the software: yum install *.rpm

This command installs the .rpm packages, all other necessary supporting software, and the Unified Manager software.



Do not attempt installation by using alternative commands (such as rpm -ivh ...). The successful installation of Unified Manager on a Red Hat Enterprise Linux or CentOS system requires that all Unified Manager files and related files are installed in a specific order into a specific directory structure that is enforced automatically by the yum install *.rpm command.

9. Disregard the email notification that is displayed immediately after the installation messages.

The email notifies the root user of an initial cron job failure, which has no adverse effect on the installation.

10. After the installation messages are complete, scroll back through the messages until you see the message in which the system displays an IP address or URL for the Unified Manager web UI, the maintenance user name (umadmin), and a default password.

The message is similar to the following:

```
OnCommand Unified Manager installed successfully.

Use a web browser and one of the following URL(s) to configure and access the Unified Manager GUI.

https://default_ip_address/ (if using IPv4)

https://[default_ip_address]/ (if using IPv6)

https://fully_qualified_domain_name/

Log in to Unified Manager in a web browser by using following details: username: umadmin password: admin
```

- 11. Record the IP address or URL, the assigned user name (umadmin), and the current password.
- 12. If you created a umadmin user account with a custom home directory prior to installing Unified Manager, then you must specify the umadmin user login shell:usermod -s /bin/maintenance-user-shell.sh umadmin

After you finish

You can access the web UI to perform the initial setup of Unified Manager, as described in the *OnCommand Unified Manager System Configuration Guide*.

Users created during Unified Manager installation

When you install Unified Manager on Red Hat Enterprise Linux or CentOS, the following users are created by Unified Manager and third-party utilities: umadmin, jboss, and mysql.

umadmin

Used to log in to Unified Manager for the first time. This user is assigned an "OnCommand Administrator" user role and is configured as the "Maintenance User" type. This user is created by Unified Manager.

· jboss

Used to run Unified Manager services related to the JBoss utility. This user is created by Unified Manager.

mysql

Used to run MySQL database queries of Unified Manager. This user is created by the MySQL third-party utility.

In addition to these users, Unified Manager also creates corresponding groups: maintenance, jboss, and mysql. The maintenance and jboss groups are created by Unified Manager, while the mysql group is created by a third-party utility.



If you created a custom home directory and defined your own umadmin user password prior to installing Unified Manager, the installation program does not recreate the maintenance group or the umadmin user.

Changing the JBoss password

You can create a new, custom JBoss password to overwrite the default password that is set during installation. This task is optional, but some sites might require this security capability to override the Unified Manager installation default setting. This operation also changes the password JBoss uses to access MySQL.

Before you begin

- You must have root user access to the Red Hat Enterprise Linux or CentOS system on which Unified Manager is installed.
- You must be able to access the NetApp-provided password.sh script in the directory /opt/netapp/essentials/bin.

Steps

- 1. Log in as root user on the system.
- 2. Stop the Unified Manager services by entering the following commands in the order shown: service ocieau stop``service ocie stop

Do not stop the associated MySQL software.

- 3. Enter the following command to begin the password change process: /opt/netapp/essentials/bin/password.sh resetJBossPassword
- 4. When prompted, enter the old JBoss password.

The default password is D11h1aMu@79%.

- 5. When prompted, enter the new JBoss password, and then enter it a second time for confirmation.
- 6. When the script completes, start the Unified Manager services by entering the following commands in the order shown: service ocie start``service ocieau start
- 7. After all of the services are started, you can log in to the Unified Manager UI.

Setting up Unified Manager for high availability

You can create a high-availability setup by using the Veritas Cluster Server (VCS). The high-availability setup provides failover capability and helps in disaster recovery.

In a high-availability setup, only one node remains active at a time. When one node fails, VCS service recognizes this event and immediately transfers control to the other node. The second node in the setup becomes active and starts providing services. The failover process is automatic.

A VCS cluster configured with the Unified Manager server consists of two nodes, with each node running the same version of the Unified Manager. All of the Unified Manager server data must be configured for access from a shared data disk.

After you install Unified Manager in VCS, you must configure Unified Manager to work in the VCS environment. You can use configuration scripts to set up Unified Manager to work in VCS environments.

Requirements for Unified Manager in VCS

Before installing Unified Manager in a Veritas Cluster Server (VCS) environment, you must ensure that the cluster nodes are properly configured to support Unified Manager.

You must ensure that the VCS configuration meets the following requirements:

- Both the cluster nodes must be running a supported operating system version.
- The same version of Unified Manager must be installed using the same path on both the cluster nodes.
- The MySQL user on both the nodes must have the same user ID and group ID.
- Native ext3, ext4 file systems, and Logical Volume Manager (LVM) must be used.
- · Unified Manager must be connected to the storage system through Fibre Channel (FC) or iSCSI.

You must also ensure that the FC link is active and that the LUNs created on the storage systems are accessible to both the cluster nodes.

- The shared data disk must have enough space (minimum 80 GB) for the Unified Manager database, reports, certificates, and script plug-in folders.
- A minimum of two network interfaces must be set up on each system: one for node-to-node communication and the other for node-to-client communication.

The name of the network interface used for node-to-client communication must be the same on both the systems.

- A separate heartbeat link must be established between the cluster nodes; otherwise, the network interface is used to communicate between the cluster nodes.
- Optional: SnapDrive for UNIX should be used to create a shared location that is accessible to both the nodes in a high availability setup.

See the *SnapDrive for UNIX Installation and Administration Guide* for information about installing and creating a shared location. You can also manage LUNs using SnapDrive or the storage system command-line interface. See the SnapDrive for UNIX compatibility matrix for more information.

• Additional RAM must be available for the SnapDrive and VCS applications.

Installing Unified Manager on VCS

For configuring high availability, you must install Unified Manager on both the cluster nodes of VCS.

Before you begin

• VCS must be installed and configured on both the nodes of the cluster.

See the instructions provided in the *Veritas Cluster Server 6.2.1 Installation Guide* for more information about installing VCS.

• You must have clear root privileges to log in to the Unified Manager server console.

About this task

You must configure both the instances of Unified Manager to use the same database and to monitor the same set of nodes.

Steps

- 1. Log in to the first node of the cluster.
- 2. Install Unified Manager on the first node.

Installing Unified Manager on Red Hat Enterprise Linux or CentOS

- Repeat Steps 1 and 2 on the second node of the cluster.
- 4. On the second instance of Unified Manager, log in as the root user to the Red Hat Enterprise Linux or CentOS server and enter the same umadmin password as you defined on the first instance of Unified Manager.passwd_umadmin

Configuring Unified Manager with VCS using configuration scripts

You can configure Unified Manager with Veritas Cluster Server (VCS) using configuration scripts.

Before you begin

- Unified Manager must be installed on both the nodes in the VCS setup.
- The XML:: LibXML module must be bundled with Perl for VCS scripts to work.
- You must have created a shared LUN with sufficient size to accommodate the source Unified Manager data.
- You must have specified the absolute mount path for the script to work.

The script will not work if you create a folder inside the mount path.

You must have downloaded the ha setup.pl script at /opt/netapp/ocum/scripts.

About this task

In the VCS setup, the node for which the virtual IP interface and mount point are active is the first node. The other node is the second node.

Steps

1. Log in to the first node of the cluster.

You must have stopped all the Unified Manager services on the second node in the high availability setup.

- 2. Add the VCS installation directory /opt/VRTSvcs/bin to the PATH environmental variable.
- 3. If you are configuring an existing Unified Manager setup, create a Unified Manager backup and generate the support bundle.
- 4. Run the ha_setup.pl script: perl ha_setup.pl --first -t vcs -g group_name -e eth_name -i cluster_ip -m net_mask -n fully_qualified_cluster_name -f mount_path -v volume_group -d disk_group -l install_dir -u user_name -p

```
perl \ha_setup.pl --first -t vcs -g umgroup -e eth0 -i 10.11.12.13 -m 255.255.255.0 -n cluster.eng.company.com -f /mnt/ocumdb -v ocumdb_SdHv -d ocumdb SdDg -l /opt/netapp/ -u admin -p wx17yz
```

- 5. Use the Veritas Operation Manager web console or VCS Cluster Manager to verify that a failover group is created, and that the Unified Manager server services, mount point, virtual IP, network interface card (NIC), and volume group are added to the cluster group.
- 6. Manually move the Unified Manager service group to the secondary node and verify that cluster failover is working.
- 7. Verify that VCS has switched over to the second node of the cluster.

You must verify that the data mount, virtual IP, volume group, and NIC are online on the second node of the cluster.

- 8. Stop Unified Manager using Veritas Operation Manager.
- 9. Run the perl ha_setup.pl --join -t vcs -f``mount_path command on the second node of the cluster so that the Unified Manager server data points to the LUN.
- 10. Verify that the Unified Manager server services are starting properly on the second node of the cluster.
- 11. Regenerate the Unified Manager certificate after running the configuration scripts to obtain the global IP address.
 - a. In the toolbar, click [1], and then click HTTPS Certificate from the Setup menu.
 - b. Click Regenerate HTTPS Certificate.

The regenerated certificate provides only the cluster IP address, not the fully qualified domain name (FQDN). You must use the global IP address to set up Unified Manager for high-availability.

12. Access the Unified Manager UI using the following: https://<FQDN of Global IP>

After you finish

You must create a shared backup location after high availability is configured. The shared location is required for containing the backups that you create before and after failover. Both the nodes in the high-availability setup must be able to access the shared location.

Unified Manager service resources for VCS configuration

You must add the cluster service resources of Unified Manager to Veritas Cluster Server (VCS). These cluster service resources are used for various purposes, such as monitoring storage systems, scheduling jobs, processing events, and monitoring all the other Unified Manager services.

The following table lists the category of all the Unified Manager services:

Category	Services
Storage resource	• vol
	* mount

Category	Services
Database resource	• mysqld
Network resource	• nic • vip
Unified Manager resource	• ocie • ocieau

Updating an existing Unified Manager setup for high availability

You can update your existing Unified Manager installation and configure your setup environment for high availability.

Before you begin

- You must have created a backup and support bundle of your existing data.
- You must have the OnCommand Administrator or Storage Administrator role.
- You must have added a second node to your cluster and installed Veritas Cluster Server (VCS) on the second node.

See the Veritas Cluster Server 6.2.1 Installation Guide.

• The newly added node must be configured to access the same shared location as that of the existing node in the high-availability setup.

Steps

- 1. Log in to the new node of the cluster.
- 2. Install Unified Manager on the node.

Installing Unified Manager on Red Hat Enterprise Linux or CentOS

- 3. Configure the Unified Manager server using configuration scripts on the existing node with data.
- 4. Initiate manual fail over to the second node.
- 5. Run the perl ha_setup.pl --join -t vcs -f``mount_path command on the second node of the cluster so that the Unified Manager server data points to the shared LUN.
- 6. If OnCommand Workflow Automation (WFA) is configured for Unified Manager, disable and then reconfigure the WFA connection.
- 7. If SnapProtect is configured with Unified Manager, reconfigure SnapProtect with a new cluster IP address and the existing storage policies.
- 8. Regenerate the custom reports and add these reports to Unified Manager with the new cluster IP address.

Upgrading Unified Manager on Red Hat Enterprise Linux or CentOS

You can upgrade Unified Manager when a new version of software is available.

Patch releases of Unified Manager software, when provided by NetApp, are installed using the same procedure as new releases.

If Unified Manager is paired with an instance of OnCommand Workflow Automation, and there are new versions of software available for both products, you must disconnect the two products and then set up a new Workflow Automation connection after performing the upgrades. If you are performing an upgrade to only one of the products, then you should log into Workflow Automation after the upgrade and verify that it is still acquiring data from Unified Manager.

Upgrading Unified Manager on Red Hat Enterprise Linux or CentOS

You can upgrade from Unified Manager version 7.3 or 9.4 to Unified Manager 9.5 by downloading and running the installation file on the Red Hat platform.

Before you begin

 The system on which you are upgrading Unified Manager must meet the system and software requirements.

Hardware system requirements

Red Hat and CentOS software and installation requirements

- Starting with Unified Manager 9.4, Red Hat Enterprise Linux 6.x is no longer supported. If you are using RHEL 6, you must upgrade your instance of RHEL to version 7.x prior to upgrading to Unified Manager 9.5.
- Starting with Unified Manager 9.5, Oracle Java is no longer supported. The correct version of OpenJDK must be installed prior to upgrading to Unified Manager 9.5.
- You must have a subscription to the Red Hat Enterprise Linux Subscription Manager.
- To avoid data loss, you must have created a backup of the Unified Manager database in case there is an issue during the upgrade. It is also recommended that you move the backup file from the /opt/netapp/data directory to an external location.
- You should have completed any running operations, because Unified Manager is unavailable during the upgrade process.

About this task



These steps contain information for systems that are configured for high availability using Veritas Operation Manager. If your system is not configured for high availability, ignore these additional steps.

Steps

- 1. Log in to the target Red Hat Enterprise Linux or CentOS server.
- 2. Download the Unified Manager bundle to the server.

Downloading Unified Manager for Red Hat or CentOS

3. Navigate to the target directory and expand the Unified Manager bundle: unzip OnCommandUnifiedManager-rhel7-9.5.zip

The required RPM modules for Unified Manager are unzipped to the target directory.

4. Confirm the presence of the listed modules: ls *.rpm

The following RPM modules are listed:

```
ocie-au-<version>.x86_64.rpm
ocie-server-<version>.x86_64.rpm
ocie-serverbase-<version>.x86_64.rpm
netapp-application-server-<version>.x86_64.rpm
netapp-platform-base-<version>.x86_64.rpm
netapp-ocum-<version>.x86_64.rpm
```

- 5. For systems that are not connected to the Internet or that are not using the RHEL repositories, perform the following steps to determine whether you are missing any required packages and download those packages:
 - a. View the list of available and unavailable packages: yum install *.rpm --assumeno

The items in the "Installing:" section are the packages that are available in the current directory, and the items in the "Installing for dependencies:" section are the packages that are missing on your system.

b. Download the missing packages on another system that has Internet access: yum install package name --downloadonly --downloaddir=.



Because the plug-in "yum-plugin-downloadonly" is not always enabled on Red Hat Enterprise Linux systems, you might need to enable the functionality to download a package without installing it: yum install yum-plugin-downloadonly

- c. Copy the missing packages from the Internet-connected system to your installation system.
- 6. If Unified Manager is configured for high availability, then using Veritas Operation Manager, stop all Unified Manager services on the first node.
- 7. Upgrade Unified Manager using the following script: upgrade.sh

This script automatically executes the RPM modules, upgrading the necessary supporting software and the Unified Manager modules that run on them. Additionally, the upgrade script checks whether there are any system configuration settings or any installed software that will conflict with the upgrade of Unified Manager. If the script identifies any issues, you must fix the issues prior to upgrading Unified Manager.



Do not attempt to upgrade by using alternative commands (such as rpm -Uvh ...). A successful upgrade requires that all Unified Manager files and related files are upgraded in a specific order to a specific directory structure that are executed and configured automatically by the script.

8. For high availability installations, stop all Unified Manager services on the second node with Veritas Operation Manager.

- 9. For high availability installations, switch the service group to the second node in the high-availability setup and upgrade Unified Manager on the second node.
- 10. After the upgrade is complete, scroll back through the messages until you see the message displaying an IP address or URL for the Unified Manager web UI, the maintenance user name (umadmin), and the default password.

The message is similar to the following:

```
OnCommand Unified Manager upgraded successfully.

Use a web browser and one of the following URLs to access the OnCommand Unified Manager GUI:

https://default_ip_address/ (if using IPv4)
https://[default_ip_address]/ (if using IPv6)
https://fully_qualified_domain_name/
```

After you finish

Enter the specified IP address or URL into a supported web browser to start the Unified Manager web UI, and then log in by using the same maintenance user name (umadmin) and password that you set earlier.

Upgrading the host OS from Red Hat Enterprise Linux 6.x to 7.x

If you previously installed Unified Manager on a Red Hat Enterprise Linux 6.x system and now need to upgrade to Red Hat Enterprise Linux 7.x, you must follow one of the procedures listed in this topic. In both cases you must create a backup of Unified Manager on the Red Hat Enterprise Linux 6.x system, and then restore the backup onto a Red Hat Enterprise Linux 7.x system.

About this task

The difference between the two options listed below is that in one case you are performing the Unified Manager restore onto a new RHEL 7.x server, and in the other case you are performing the restore operation onto the same server.

Because this task requires that you create a backup of Unified Manager on the Red Hat Enterprise Linux 6.x system, you should create the backup only when you are prepared to complete the entire upgrade process so that Unified Manager is offline for the shortest period of time. Gaps in collected data will appear in the Unified Manager UI for the period of time during which the Red Hat Enterprise Linux 6.x system is shut down and before the new Red Hat Enterprise Linux 7.x is started.

See the *Unified Manager Online Help* if you need to review detailed instructions for the backup and restore processes.

Upgrading the host OS using a new server

Follow these steps if you have a spare system on which you can install RHEL 7.x software so that you can perform the Unified Manager restore on that system while the RHEL 6.x system is still available.

1. Install and configure a new server with Red Hat Enterprise Linux 7.x software.

Red Hat software and installation requirements

2. On the Red Hat Enterprise Linux 7.x system, install the same version of Unified Manager software that you have on the existing Red Hat Enterprise Linux 6.x system.

Installing Unified Manager on Red Hat Enterprise Linux

Do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.

- 3. On the Red Hat Enterprise Linux 6.x system, from the Administration menu in the web UI, create a Unified Manager backup and then copy the backup file to an external location.
- 4. On the Red Hat Enterprise Linux 6.x system, shut down Unified Manager.
- 5. On the Red Hat Enterprise Linux 7.x system, copy the backup file from the external location to /data/ocum-backup/, and then enter the following command to restore the Unified Manager database from the backup file:um backup restore -f /opt/netapp/data/ocumbackup/<backup_file_name>
- 6. Enter the IP address or URL into a supported web browser to start the Unified Manager web UI, and then log in to the system.

Once you have verified that the system is operating properly you can remove Unified Manager from the Red Hat Enterprise Linux 6.x system.

Upgrading the host OS on the same server

Follow these steps if you do not have a spare system on which you can install RHEL 7.x software.

- 1. From the Administration menu in the web UI, create a Unified Manager backup and then copy the backup file to an external location.
- 2. Remove the Red Hat Enterprise Linux 6.x image from the system and completely wipe the system.
- 3. Install and configure Red Hat Enterprise Linux 7.x software on the same system.

Red Hat software and installation requirements

4. On the Red Hat Enterprise Linux 7.x system, install the same version of Unified Manager software that you had on the Red Hat Enterprise Linux 6.x system.

Installing Unified Manager on Red Hat Enterprise Linux

Do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.

- 5. Copy the backup file from the external location to /data/ocum-backup/, and then enter the following command to restore the Unified Manager database from the backup file:um backup restore -f /opt/netapp/data/ocum-backup/

 sackup file name>
- 6. Enter the IP address or URL into a supported web browser to start the Unified Manager web UI, and then log in to the system.

Upgrading third-party products on Linux

You can upgrade third-party products, such as JRE and MySQL, on Unified Manager

when installed on Linux systems.

The companies that develop these third-party products report security vulnerabilities on a regular basis. You can upgrade to newer versions of this software at your own schedule.

Upgrading JRE on Linux

You can upgrade to a newer version of Java Runtime Environment (JRE) on the Linux server on which Unified Manager is installed to obtain fixes for security vulnerabilities.

Before you begin

You must have root privileges for the Linux system on which Unified Manager is installed.

Steps

- 1. Log in as a root user on the Unified Manager host machine.
- 2. Download the appropriate version of Java (64-bit) to the target system.
- 3. Stop the Unified Manager services: service ocieau stop``service ocie stop
- 4. Install the latest JRE on the system.
- 5. Start the Unified Manager services: service ocie start ``service ocieau start

Upgrading MySQL on Linux

You can upgrade to a newer version of MySQL on the Linux server on which Unified Manager is installed to obtain fixes for security vulnerabilities.

Before you begin

You must have root privileges for the Linux system on which Unified Manager is installed.

About this task

You can only upgrade to minor updates of MySQL 5.7, for example, 5.7.1 to 5.7.2 . You cannot upgrade to major versions of MySQL, for example, version 5.8.

- 1. Log in as a root user on the Unified Manager host machine.
- 2. Download the latest MySQL Community Server .rpm bundle on the target system.
- 3. Untar the bundle to a directory on the target system.
- 4. You will get multiple .rpm packages in the directory after untarring the bundle, but Unified Manager only needs the following rpm packages:
 - mysql-community-client-5.7.x
 - mysql-community-libs-5.7.x
 - mysql-community-server-5.7.x
 - mysql-community-common-5.7.x
 - ° mysql-community-libs-compat-5.7.x

Delete all other .rpm packages. Installing all packages in an rpm bundle will not cause any problems.

- 5. Stop the Unified Manager service and the associated MySQL software in the order shown: service ocieau stopservice ocie stopservice mysqld stop
- 6. Invoke the upgrade of MySQL by using the following command: yum install *.rpm
 - *.rpm refers to the .rpm packages in the directory where you downloaded the newer version of MySQL.
- 7. Start Unified Manager in the order shown: service mysqld startservice ocie startservice ocieau start

Restarting Unified Manager in Red Hat Enterprise Linux or CentOS

You might have to restart Unified Manager after making configuration changes.

Before you begin

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

Steps

- 1. Log in as root user to the server on which you want to restart the Unified Manager service.
- 2. Stop the Unified Manager service and the associated MySQL software in the order shown: service ocieau stopservice ocie stopservice mysqld stop
 - When installed in a high-availability setup, stop the Unified Manager service by using either VCS Operations Manager or VCS commands.
- 3. Start Unified Manager in the order shown: service mysqld startservice ocie startservice ocieau start

When installed in a high-availability setup, start Unified Manager service by using either VCS Operations Manager or VCS commands.

Removing Unified Manager from the Red Hat Enterprise Linux or CentOS host

If you need to remove Unified Manager from the Red Hat Enterprise Linux or CentOS host, you can stop and uninstall Unified Manager with a single command.

Before you begin

- You must have root user access to the server from which you want to remove Unified Manager.
- Security-Enhanced Linux (SELinux) must be disabled on the Red Hat machine. Change the SELinux runtime mode to "Permissive" by using the setenforce 0 command.
- All clusters (data sources) must be removed from the Unified Manager server before removing the software.
- The Unified Manager server must not have an active connection to an external data provider such as Graphite.

If it does, you must delete the connection using the Unified Managermaintenance console.

About this task

These steps contain information for systems that are configured for high availability using Veritas Operation Manager. If your system is not configured for high availability, ignore these additional steps.

Steps

- 1. Log in as root user to the cluster node owning the cluster resources on which you want to remove Unified Manager.
- 2. Stop all Unified Manager services using VCS Operations Manager or VCS commands.
- 3. Stop and remove Unified Manager from the server: rpm -e netapp-ocum ocie-au ocie-server netapp-platform-base netapp-application-server ocie-serverbase

This step removes all the associated NetApp RPM packages. It does not remove the prerequisite software modules, such as Java, MySQL, and p7zip.

- 4. Switch to the other node by using the VCS Operations Manager.
- 5. Log in to the second node of the cluster.
- 6. Stop all the services, and then and remove Unified Manager from the second node: rpm -e netappocum ocie-au ocie-server netapp-platform-base netapp-application-server ocieserverbase
- 7. Prevent the service group from using VCS Operations Manager or VCS commands.
- 8. If appropriate, remove the supporting software modules, such as Java, MySQL, and p7zip: rpm -e p7zip mysql-community-client mysql-community-server mysql-community-common mysql-community-libs java-x.y

Results

After this operation is complete, the software is removed; however, MySQL data is not deleted. All the data from the <code>/opt/netapp/data directory</code> is moved to the <code>/opt/netapp/data/BACKUP</code> folder after uninstallation.

Removing the custom umadmin user and maintenance group

If you created a custom home directory to define your own umadmin user and maintenance account prior to installing Unified Manager, you should remove these items after you have uninstalled Unified Manager.

About this task

The standard Unified Manager uninstallation does not remove a custom-defined umadmin user and maintenance account. You must delete these items manually.

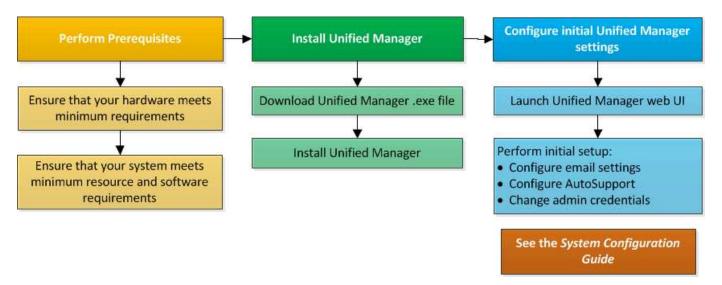
- 1. Log in as the root user to the Red Hat Enterprise Linux server.
- 2. Delete the umadmin user:userdel umadmin

Installing, upgrading, and removing Unified Manager software on Windows

On Windows systems, you can install Unified Manager software, upgrade to a newer version of software, or remove the Unified Manager application.

Overview of the installation process on Windows

The installation workflow describes the tasks that you must perform before you can use Unified Manager.



Installing Unified Manager on Windows

It is important that you understand the sequence of steps to download and install Unified Manager on Windows. Before you install Unified Manager on Windows, you can decide if you want to configure Unified Manager for high availability.

Installing Unified Manager on a Windows system

You can install Unified Manager on Windows to monitor and troubleshoot data storage capacity, availability, performance, and protection issues.

Before you begin

• The system on which you plan to install Unified Manager must meet the system and software requirements.

Hardware system requirements

Windows software and installation requirements



Starting with Unified Manager 9.5, OpenJDK is provided in the Unified Manager installation package and installed automatically. Oracle Java is not supported starting with Unified Manager 9.5.

- · You must have Windows administrator privileges.
- · You must have a supported web browser.
- The Unified Manager maintenance user password must be between 8 and 20 characters, must contain upper-case letters or lower-case letters, numerals, and special characters.
- The following special characters are not allowed in the password string for the maintenance user or for the MySQL root user: "'`%, = & < > | ^ \ / ()[];

The following special characters are allowed: ~! @ # \$ * - ? . : + { }

Steps

- 1. Log in to Windows using the default local administrator account.
- Log in to the NetApp Support Site, and locate the Download page for installing Unified Manager on the Windows platform.

NetApp Downloads: Software

- Download the Unified Manager Windows installation file from the NetApp Support Site to a target directory in the Windows system.
- 4. Navigate to the directory where the installation file is located.
- Right-click and run the Unified Manager installer executable (.exe) file as an administrator.

Unified Manager detects missing or pre-installed third-party packages and lists them. If the required third-party packages are not installed in the system, Unified Manager installs them as part of the installation.

- 6. Click Next.
- 7. Enter the user name and password to create the maintenance user.
- 8. In the Database Connection wizard, enter the MySQL root password.
- Click Change to specify a new location for the Unified Manager installation directory and MySQL data directory.

If you do not change the installation directory, Unified Manager is installed in the default installation directory.

- 10. Click Next.
- 11. In the Ready to Install Shield wizard, click Install.
- 12. After the installation is complete, click Finish.

Results

The installation creates multiple directories:

Installation directory

This is the root directory for Unified Manager, which you specified during installation. Example:

C:\Program Files\NetApp\

· MySQL data directory

This is the directory where the MySQL databases are stored, which you specified during installation. **Example:** C:\ProgramData\MySQL\MySQLServerData\

· Java directory

This is the directory where OpenJDK will be installed. Example: C:\Program Files\NetApp\JDK\

Unified Manager application data directory (appDataDir)

This is the directory where all the application-generated data is stored. This includes logs, support bundles, backup, and all other additional data. Example: C:\ProgramData\NetApp\OnCommandAppData\

After you finish

You can access the web UI to perform the initial setup of Unified Manager, as described in Configuring Unified Manager.

Performing an unattended installation of Unified Manager

You can install Unified Manager without user intervention by using the command-line interface. You can complete the unattended installation by passing the parameters in key-value pairs.

- 1. Log in to the Windows command-line interface by using the default local administrator account.
- 2. Navigate to the location where you want to install Unified Manager, and choose one of the following options:

Option	Instructions
If third-party packages are pre-installed	OnCommandUnifiedManager-x.y.exe /V"MYSQL_PASSWORD=mysql_password INSTALLDIR=\"Installation directory\" MYSQL_DATA_DIR=\"MySQL data directory\" MAINTENANCE_PASSWORD=maintenance_passw ord MAINTENANCE_USERNAME=maintenance_usern ame /qn /l*v CompletePathForLogFile"
	Example:
	OnCommandUnifiedManager.exe /s /v"MYSQL_PASSWORD=netapp21! INSTALLDIR=\"C:\Program Files\NetApp\" MYSQL_DATA_DIR=\"C:\ProgramData\MYSQL\ MySQLServer\" MAINTENANCE_PASSWORD=* MAINTENANCE_USERNAME=admin /qn /l*v C:\install.log"
If third-party packages are not installed	OnCommandUnifiedManager-x.y.exe /V"MYSQL_PASSWORD=mysql_password INSTALLDIR=\"Installation directory\" MYSQL_DATA_DIR=\"MySQL data directory\" MAINTENANCE_PASSWORD=maintenance_passw ord MAINTENANCE_USERNAME=maintenance_usern ame /qr /l*v CompletePathForLogFile"
	Example:
	OnCommandUnifiedManager.exe /s /v"MYSQL_PASSWORD=netapp21! INSTALLDIR=\"C:\Program Files\NetApp\" MYSQL_DATA_DIR=\"C:\ProgramData\MYSQL\ MYSQLServer\" MAINTENANCE_PASSWORD=* MAINTENANCE_USERNAME=admin /qr /l*v C:\install.log"

The / qr option enables quiet mode with a reduced user interface. A basic user interface is displayed, which shows the installation progress. You are not prompted for inputs. If third-party packages such as JRE, MySQL, and 7zip are not pre-installed, you must use the / qr option. Installation fails if the / qn option is used on a server where third-party packages are not installed.

The /qn option enables quiet mode with no user interface. No user interface or details are displayed during installation. You must not use the /qn option when third-party packages are not installed.

3. Log in to the Unified Manager web user interface by using the following URL:

Setting up Unified Manager in a failover clustering environment

You can configure high availability for Unified Manager using failover clustering. The high-availability setup provides failover capability.

In this setup, only one node owns all the cluster resources. When one node goes down or any of the configured services fail to come online, the failover cluster service recognizes this event and immediately transfers control to the other node. The second node in the setup becomes active and starts providing services. The failover process is automatic and you do not have to perform any actions.

A failover cluster configured with the Unified Manager server consists of two nodes, each node running the same version of the Unified Manager server. All of the Unified Manager server data must be configured for access from a shared data disk.

Requirements for Unified Manager in a failover clustering environment

Before installing Unified Manager in a failover clustering environment, you must ensure that the cluster nodes are properly configured to support Unified Manager.

You must ensure that the failover cluster configuration meets the following requirements:

- Both the cluster nodes must be running the same version of Microsoft Windows Server.
- The same version of Unified Manager must be installed using the same path on both the cluster nodes.
- Failover clustering must be installed and enabled on both the nodes.

See Microsoft documentation for instructions.

- You must have used Fibre Channel switched fabric or iSCSI-based storage for creating shared data disk as the storage back-end.
- Optional: Using SnapDrive for Windows, a shared location must be created that is accessible to both the nodes in the high-availability setup.

See the *SnapDrive for Windows Installation Guide* for information about installing and creating a shared location.

You can also manage LUNs using the storage system command-line interface. See the SnapDrive for Windows compatibility matrix for more information.

- You must have the Perl installed with XML::LibXML and File::chdir modules for scripts to work.
- There must be only two nodes in the cluster setup.
- The "node and disk majority" quorum type must be used for failover clustering.
- You must have configured a shared IP address with a corresponding FQDN to be used as the cluster global IP address to access Unified Manager.
- The password for Unified Manager maintenance user on both the nodes must be same.
- · You must have used only IPv4 IP address.

Installing Unified Manager on MSCS

For configuring high availability, you must install Unified Manager on both the Microsoft Cluster Server (MSCS) cluster nodes.

Steps

- 1. Log in as the domain user on both the nodes of the cluster.
- 2. Set up high availability by choosing one of the following options:

If you want to	Then do this
Configure high availability on an existing Unified Manager installation	Add another server to be paired with the existing server:
	Upgrade the existing Unified Manager server to the latest software version.
	 b. Create a backup of the existing Unified Manager installation, and store the backup to a mounted LUN.
	c. Install Unified Manager on the second node.
	Installing Unified Manager on a Windows system
	d. Restore the backup of the existing Unified Manager installation onto the second node.
Configure high availability on a new Unified Manager installation	Install Unified Manager on both the nodes. Installing Unified Manager on a Windows system

Configuring Unified Manager server with MSCS using configuration scripts

After installing Unified Manager on both cluster nodes, you can configure Unified Manager with Failover Cluster Manager using configuration scripts.

Before you begin

You must have created a shared LUN that is of a sufficient size to accommodate the source Unified Manager data.

- 1. Log in to the first node of the cluster.
- 2. Create a role in Windows 2012 or Windows 2016 using Failover Cluster Manager:
 - a. Launch Failover Cluster Manager.
 - b. Create the empty role by clicking Roles > Create Empty Role.
 - c. Add the global IP address to the role by right-clicking Role > Add Resources > More Resources > IP address.



Both nodes must be able to ping this IP address because Unified Manager is launched using this IP address after high availability is configured.

- d. Add the data disk to the role by right-clicking Role > Add Storage.
- 3. Run the ha_setup.pl script on the first node: perl ha_setup.pl --first -t mscs -g group_name -i ip address -n fully_qualified_domain_cluster_name -f shared_location_path -k data_disk -u user_name -p password

```
C:\Program Files\NetApp\ocum\bin>perl .\ha_setup.pl --first -t mscs -g umgroup
-i "IP Address" -n spr38457002.eng.company.com -k "Cluster Disk 2" -f E:\ -u
admin -p wx17yz
```

The script is available at Install Dir\NetApp\ocum\bin.

- You can obtain the value of the -g, -k, and -i options using the cluster res command.
- ∘ The -n option must be the FQDN of the global IP address that can be pinged from both nodes.
- 4. Verify that the Unified Manager server services, data disk, and cluster IP address are added to the cluster group by using the Failover Cluster Manager web console.
- 5. Stop all Unified Manager server services (MySQL, ocie, and ocieau) by using the services.msc command.
- 6. Switch the service group to the second node in Failover Cluster Manager.
- 7. Run the command perl ha_setup.pl --join -t mscs -f``shared_location_path on the second node of the cluster to point to the Unified Manager server data to the LUN.

```
perl ha setup.pl --join -t mscs -f E:\
```

- Bring all the Unified Manager services online using Failover Cluster Manager.
- 9. Manually switch to the other node of the Microsoft Cluster Server.
- 10. Verify that the Unified Manager server services are starting properly on the other node of the cluster.
- 11. Regenerate the Unified Manager certificate after running configuration scripts to obtain the global IP address.
 - a. In the toolbar, click , and then click HTTPS Certificate from the Setup menu.
 - b. Click Regenerate HTTPS Certificate.

The regenerated certificate provides the cluster IP address, not the fully qualified domain name (FQDN). You must use the global IP address to set up Unified Manager for high-availability.

12. Access the Unified Manager UI using the following: https://<FQDN of Global IP>

After you finish

You must create a shared backup location after high availability is configured. The shared location is required for containing the backups before and after failover. Both nodes in the high-availability setup must be able to access the shared location.

Upgrading Unified Manager on Windows

You can upgrade Unified Manager 7.3 or 9.4 to Unified Manager 9.5 by downloading and running the installation file on the Windows platform.

Before you begin

• The system on which you are upgrading Unified Manager must meet the system and software requirements.

Hardware system requirements

Windows software and installation requirements



Starting with Unified Manager 9.5, OpenJDK is provided in the Unified Manager installation package and installed automatically. Oracle Java is not supported starting with Unified Manager 9.5.



Starting with Unified Manager 9.4, Microsoft .NET 4.5.2 or greater is required. Make sure you have the correct version of .NET installed before starting the upgrade.

- · You must have Windows administrator privileges.
- You must have valid credentials to log in to the NetApp Support Site.
- To avoid data loss, you must have created a backup of the Unified Manager machine in case there is an issue during the upgrade.
- · You must have adequate disk space available to perform the upgrade.

The available space on the installation drive must be 2.5 GB larger than the size of the data directory. The upgrade will stop and display an error message indicating the amount of space to be added if there is not enough free space.

About this task

During the upgrade process, Unified Manager is unavailable. You should complete any running operations before upgrading Unified Manager.

If Unified Manager is paired with an instance of OnCommand Workflow Automation, and there are new versions of software available for both products, you must disconnect the two products and then set up a new Workflow Automation connection after performing the upgrades. If you are performing an upgrade to only one of the products, then you should log into Workflow Automation after the upgrade and verify that it is still acquiring data from Unified Manager.

Steps

1. Log in to the NetApp Support Site, and locate the Download page for installing Unified Manager on the Windows platform.

NetApp Downloads: Software

- 2. Download the Unified Manager Windows installation file to a target directory in the Windows system.
- 3. If Unified Manager is configured for high availability, stop all the Unified Manager services on the first node

by using Microsoft Cluster Server, and then start the MySQL service from services.msc.

4. Right-click and run the Unified Manager installer executable (.exe) file as an administrator.

Unified Manager prompts you with the following message:

```
This setup will perform an upgrade of 'OnCommand Unified Manager'. Do you want to continue?
```

- 5. Click Yes, and then click Next.
- Enter the MySQL root password that was set during installation, and click Next.
- 7. After the upgrade is successful, if the system is configured for high availability, start all the Unified Manager services from the Failover Cluster Manager and follow the remaining tasks.
- 8. From the command prompt, run the ha_setup.pl script to configure the new services in the failover cluster and the files that are present in the shared location.

```
C:\Program Files\NetApp\ocum\bin> perl .\ha_setup.pl --upgrade --first -t mscs
-g kjaggrp -i "New IP Address1" -n scs8003.englab.company.com -k "Cluster Disk
2" -f E:\ -u user -p userpass
```

- 9. Stop all the Unified Manager services (ocie, ocieau, and MySQL) in the first node by using Microsoft Cluster Server.
- 10. Start the MySQL service on the second node from services.msc.
- 11. Switch the service group to the second node in the high-availability setup.
- 12. Upgrade Unified Manager on the second node.
- 13. At the command prompt, enter Y to continue, or enter any other character to abort.

The upgrade and restart processes of the Unified Manager services can take several minutes to complete.

- 14. Start all the Unified Manager services on both the nodes using Microsoft Cluster Server.
- 15. From the command prompt, run the ha setup.pl script with the --upgrade option.

```
perl ha setup.pl --upgrade --join -t mscs -f E:\
```

16. Log in to the Unified Manager web UI, and verify the version number.

After you finish



To perform a silent upgrade of Unified Manager, run the following command:

OnCommandUnifiedManager-9.5.exe /s /v"MYSQL_PASSWORD=netapp21! /qn /l*vC:\install.log

Upgrading third-party products on Windows

You can upgrade third-party products, such as JRE and MySQL, on Unified Manager when installed on Windows systems.

The companies that develop these third-party products report security vulnerabilities on a regular basis. You can upgrade to newer versions of this software at your own schedule.

Upgrading JRE on Windows

You can upgrade to a newer version of Java Runtime Environment (JRE) on the Windows server on which Unified Manager is installed to obtain fixes for security vulnerabilities.

Before you begin

You must have Windows admin privileges for the system on which Unified Manager is installed.

Steps

- 1. Log in as the admin user on the Unified Manager host machine.
- 2. Download the appropriate version of Java (64-bit) from the JDK site to the target system.

For example, download openjdk-11_windows-x64_bin.zip from http://jdk.java.net/11/.

- 3. Use the Windows Services console to stop the following Unified Manager services:
 - NetApp OCIE Acquisition Unit (Ocie-au)
 - NetApp OnCommand Application Server (Oncommandsvc)
- 4. Expand the zip file.
- 5. Copy the directories and files from the resulting jdk directory (for example, jdk-11.0.1 to the location where Java is installed. Example: C:\Program Files\NetApp\JDK\
- 6. Start the Unified Manager services by using the Windows Services console:
 - NetApp OnCommand Application Server (Oncommandsvc)
 - NetApp OCIE Acquisition Unit (Ocie-au)

Upgrading MySQL on Windows

You can upgrade to a newer version of MySQL on the Windows server on which Unified Manager is installed to obtain fixes for security vulnerabilities.

Before you begin

- You must have Windows admin privileges for the system on which Unified Manager is installed.
- You must have the password for the MySQL root user.

- 1. Log in as the admin user on the Unified Manager host machine.
- 2. Download the appropriate version of MySQL to the target system.
- 3. Use the Windows Services console to stop the following Unified Manager services:
 - NetApp OCIE Acquisition Unit (Ocie-au)
 - NetApp OnCommand Application Server (Oncommandsvc)
 - MYSQL

- 4. Click the .msi package to invoke the upgrade of MySQL and follow the instructions on the screen to complete the upgrade.
- 5. Start the Unified Manager services by using the Windows Services console:
 - MYSQL
 - NetApp OnCommand Application Server (Oncommandsvc)
 - · NetApp OCIE Acquisition Unit (Ocie-au)

Restarting Unified Manager on Windows

You might have to restart Unified Manager after making configuration changes.

Before you begin

You must have Windows administrator privileges.

Steps

- 1. Log in to Windows using the default local administrator account.
- 2. Stop the Unified Manager services:

From the	Stop the services in following order
Command line	a. sc stop ocie-aub. sc stop Oncommandsvc
Microsoft Service Manager	a. NetApp OCIE Acquisition Unit (Ocie-au)b. NetApp OnCommand Application Server (Oncommandsvc)

When installed in a high-availability setup, stop the Unified Manager service by using either Microsoft Service Manager or the command line.

3. Start the Unified Manager services:

From the	Start the services in following order
Command line	a. sc start Oncommandsvcb. sc start ocie-au
Microsoft Service Manager	a. NetApp OnCommand Application Server (Oncommandsvc)b. NetApp OCIE Acquisition Unit (Ocie-au)

When installed in a high-availability setup, start Unified Manager service by using either Microsoft Service Manager or the command line.

Uninstalling Unified Manager from Windows

You can uninstall Unified Manager from Windows by using the Programs and Features wizard, or by performing an unattended uninstallation from the command-line interface.

Before you begin

- · You must have Windows administrator privileges.
- All clusters (data sources) must be removed from the Unified Manager server before uninstalling the software.
- The Unified Manager server must not have an active connection to an external data provider such as Graphite.

If it does, you must delete the connection using the Unified Managermaintenance console.

Steps

- 1. When installed in a high-availability setup, remove the HA service group resources and delete the HA service group before uninstalling Unified Manager from both nodes.
- 2. Uninstall Unified Manager by choosing one of the following options:

To uninstall Unified Manager from the	Then
Programs and Features wizard	a. Navigate to Control Panel > Program and Features.
	b. Select OnCommand Unified Manager, and click Uninstall.
Command line	a. Log in to the Windows command line using administrator privileges.
	b. Navigate to the OnCommand Unified Manager directory, and run the following command: msiexec /x {A78760DB-7EC0-4305- 97DB-E4A89CDFF4E1} /qn /1*v %systemdrive%\UmUnInstall.log

If User Account Control (UAC) is enabled on the server, and you are logged in as a domain user, you must use the command-line uninstallation method.

Unified Manager is uninstalled from your system.

- 3. Uninstall the following third-party packages and data that are not removed during the Unified Manager uninstallation:
 - Third-party packages: JRE, MySQL, Microsoft Visual C++ 2015 Redistributable, and 7zip
 - MySQL application data generated by Unified Manager
 - Application logs and contents of application data directory

Perform configuration and administrative tasks

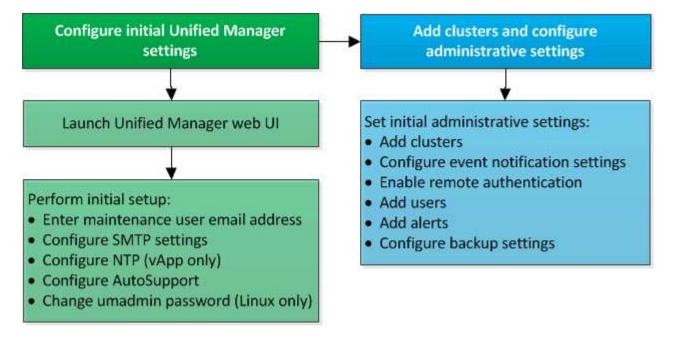
Configuring Unified Manager

After installing Unified Manager you must complete the initial setup (also called the first experience wizard) to access the web UI. Then you can perform additional configuration tasks, such as adding clusters, configuring remote authentication, adding users, and adding alerts.

Some of the procedures described in this manual are required to complete the initial setup of your Unified Manager instance. Other procedures are recommended configuration settings that are helpful to set up on your new instance, or that are good to know about before you start the regular monitoring of your ONTAP systems.

Overview of the configuration sequence

The configuration workflow describes the tasks that you must perform before you can use Unified Manager.



Accessing the Unified Manager web UI

After you have installed Unified Manager, you can access the web UI to set up Unified Manager so that you can begin monitoring your ONTAP systems.

Before you begin

- If this is the first time you are accessing the web UI, you must log in as the maintenance user (or umadmin user for Linux installations).
- If you plan to allow users to access Unified Manager using the short name instead of using the fully
 qualified domain name (FQDN) or IP address, then your network configuration has to resolve this short
 name to a valid FQDN.

• If the server uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate for server authentication.

Steps

 Start the Unified Manager web UI from your browser by using the URL displayed at the end of the installation. The URL is the IP address or fully qualified domain name (FQDN) of the Unified Manager server.

The link is in the following format: https://URL.

2. Log in to the Unified Manager web UI using your maintenance user credentials.

Performing the initial setup of the Unified Manager web UI

To use Unified Manager, you must first configure the initial setup options, including the NTP server, the maintenance user email address, and the SMTP server host name and options.

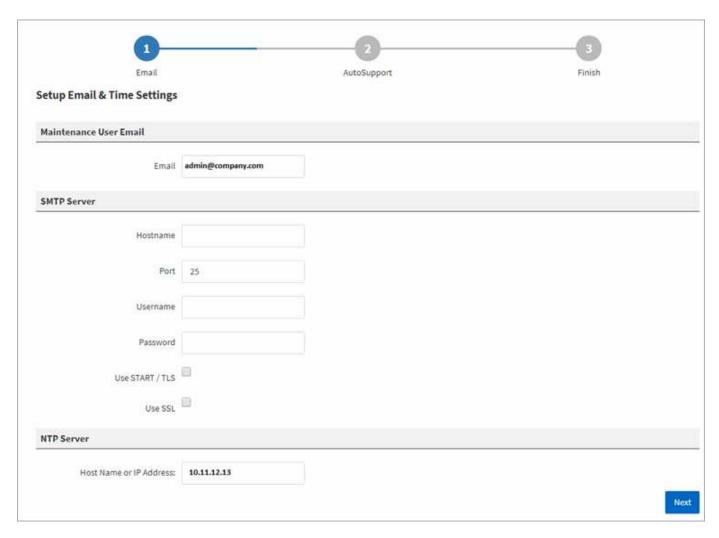
Before you begin

You must have performed the following operations:

- Launched the Unified Manager web UI using the URL provided after installation
- Logged in using the maintenance user name and password (umadmin user for Linux installations) created during installation

About this task

The OnCommand Unified Manager Initial Setup page appears only when you first access the web UI. The page below is from an installation on VMware.



If you want to change any of these options later, you can use the Administration options, which are accessible by clicking the from the Unified Manager toolbar.

Steps

- In the OnCommand Unified Manager Initial Setup window, enter the maintenance user email address, the SMTP server host name and any additional SMTP options, and the NTP server (VMware installations only). Then click Next.
- 2. In the AutoSupport page click Agree and Continue to enable AutoSupport.

If you need to designate a proxy to provide internet access in order to send AutoSupport content to support, or if you want to disable AutoSupport, use the Administration options.

On Red Hat and CentOS systems you can choose to change the umadmin user password from the default "admin" string to a personalized string.

Results

The Initial Setup window closes and the Unified Manager web UI is displayed. The Configuration/Cluster Data Sources page appears so that you can add clusters to your system.

Adding clusters

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

Before you begin

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

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

The cluster-management IP address must be the cluster-management LIF of the administrative storage virtual machine (SVM). If you use a node-management LIF, the operation fails.

Data ONTAP administrator user name and password

This account must have the admin role with Application access set to ontapi, ssh, and http.

 Type of protocol (HTTP or HTTPS) that can be configured on the cluster and the port number used to connect to the cluster



You can add clusters which are behind a NAT/firewall by using the Unified Manager NAT IP address. Any connected Workflow Automation or SnapProtect systems must also be behind the NAT/firewall, and SnapProtect API calls must use the NAT IP address to identify the cluster.

• The Unified Manager FQDN must be able to ping the ONTAP system.

You can verify this by using the following ONTAP command: ping -node node_name -destination Unified Manager FQDN.

• You must have adequate space on the Unified Manager server. You are prevented from adding a cluster to the server when greater than 90% of space in the database directory is already consumed.

About this task

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

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

- 1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
- 2. On the Configuration/Cluster Data Sources page, click Add.

3. In the **Add Cluster** dialog box, specify the required 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 and port 443 are selected.

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

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

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

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

Results

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



Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time. Additionally, if you restart Unified Manager during the data continuity collection period, the collection will be halted and you will see gaps in the performance charts for the missing timeframe.

If you receive an error message that you cannot add the cluster, check to see if the following issues exist:



- If the clocks on the two systems are not synchronized and the Unified Manager HTTPS certificate start date is later than the date on the cluster. You must ensure that the clocks are synchronized using NTP or a similar service.
- If the cluster has reached the maximum number of EMS notification destinations the Unified Manager address cannot be added. By default only 20 EMS notification destinations can be defined on the cluster.

Configuring Unified Manager to send alert notifications

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

Before you begin

You must have the OnCommand Administrator role.

About this task

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

Steps

1. Configure event notification settings

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

2. Enable remote authentication

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

3. Add authentication servers

You can add authentication servers so that remote users within the authentication server can access Unified Manager.

4. Add users

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

5. Add alerts

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

Configuring event notification settings

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

Before you begin

You must have the following information:

· Email address from which the alert notification is sent

The email address appears in the "From" field in sent alert notifications. If the email cannot be delivered for any reason, this email address is also used as the recipient for undeliverable mail.

- · SMTP server host name, and the user name and password to access the server
- SNMP version, trap destination host IP address, outbound trap port, and the community to configure the SNMP trap

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

Notes:

- If the From Address is pre-filled with the address "OnCommand@localhost.com", you should change it to a real, working email address to make sure that all email notifications are delivered successfully.
- If the host name of the SMTP server cannot be resolved, you can specify the IP address (IPv4 or IPv6)
 of the SMTP server instead of the host name.

Enabling remote authentication

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

Before you begin

You must have the OnCommand Administrator role.



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

About this task

You can enable remote authentication using either Open LDAP or Active Directory. If remote authentication is disabled, remote users cannot access Unified Manager.

Remote authentication is supported over LDAP and LDAPS (Secure LDAP). Unified Manager uses 389 as the default port for non-secure communication, and 636 as the default port for secure communication.



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

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

For Authentication type	Enter the following information
Active Directory	Authentication server administrator name in one of following formats:
	° domainname\username
	° username@domainname
	° Bind Distinguished Name (using the appropriate LDAP notation)
	Administrator password
	 Base distinguished name (using the appropriate LDAP notation)
Open LDAP	 Bind distinguished name (in the appropriate LDAP notation)
	Bind password
	Base distinguished name

If the authentication of an Active Directory user takes a long time or times out, the authentication server is probably taking a long time to respond. Disabling support for nested groups in Unified Manager might reduce the authentication time.

If you select the Use Secure Connection option for the authentication server, then Unified Manager communicates with the authentication server using the Secure Sockets Layer (SSL) protocol.

- 4. Add authentication servers, and test the authentication.
- 5. Click Save and Close.

Disabling nested groups from remote authentication

If you have remote authentication enabled, you can disable nested group authentication so that only individual users, and not group members, can remotely authenticate to Unified Manager. You can disable nested groups when you want to improve Active Directory authentication response time.

Before you begin

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

About this task

Disabling support for nested groups in Unified Manager might reduce the authentication time. If nested group support is disabled, and if a remote group is added to Unified Manager, individual users must be members of the remote group to authenticate to Unified Manager.

Steps

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

- 2. In the Setup/Authentication page, check the Disable Nested Group Lookup box.
- 3. Click Save.

Adding authentication servers

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

Before you begin

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

About this task

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

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, click Management Server > Authentication.
- 3. Enable or disable the Use secure connection authentication option:

If you want to	Then do this	
Enable it	 a. In Enable remote authentication checkbox, select the Use Secure Connection option. b. In the Authentication Servers area, click Add. c. In the Add Authentication Server dialog box, enter the authentication name or IP address (IPv4 or IPv6) of the server. d. In the Authorize Host dialog box, click View Certificate. e. In the View Certificate dialog box, verify the certificate information, and then click Close. f. In the Authorize Host dialog box, click Yes. 	
	i	When you enable the Use Secure Connection authentication option, Unified Manager communicates with the authentication server and displays the certificate. Unified Manager uses 636 as default port for secure communication and port number 389 for non-secure communication.
Disable it	 a. In the Enable remote authentication checkbox, clear the Use Secure Connection option. b. In the Authentication Servers area, click Add. c. In the Add Authentication Server dialog box, 	
	specify either the host name or IP address (IPv4 or IPv6) of the server, and the port details.	
	d. Click Add	d.

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

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

Testing the configuration of authentication servers

You can validate the configuration of your authentication servers to ensure that the management server is able to communicate with them. You can validate the configuration by searching for a remote user or remote group from your authentication servers, and authenticating them using the configured settings.

Before you begin

- You must have enabled remote authentication, and configured your authentication service so that the Unified Manager server can authenticate the remote user or remote group.
- You must have added your authentication servers so that the management server can search for the remote user or remote group from these servers and authenticate them.
- · You must have the OnCommand Administrator role.

About this task

If the authentication service is set to Active Directory, and if you are validating the authentication of remote users who belong to the primary group of the authentication server, information about the primary group is not displayed in the authentication results.

Steps

- 1. In the toolbar, click 🚺, and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, click Test Authentication.
- 3. In the **Test User** dialog box, specify the user name and password of the remote user or the user name of the remote group, and then click **Test**.

If you are authenticating a remote group, you must not enter the password.

Adding users

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

Before you begin

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

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

About this task

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

Steps

1. In the toolbar, click 🚺, and then click **Users** in the left Management menu.

- On the Management/Users page, click Add.
- In the Add User dialog box, select the type of user that you want to add, and enter the required information.

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

4. Click Add.

Adding alerts

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

Before you begin

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

About this task

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

Steps

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

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

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

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



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

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



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

You can also choose to notify users through SNMP traps.

Click Save.

Example of adding an alert

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

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

Perform the following steps in the Add Alert dialog box:

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

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

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

EMS events that are added automatically to Unified Manager

When using Unified Manager 9.4 or greater software, the following ONTAP EMS events are added automatically to Unified Manager. These events will be generated when triggered on any cluster that Unified Manager is monitoring.

The following EMS events are available when monitoring clusters running ONTAP 9.5 or greater software:

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
Object-store Access Denied for Aggregate Relocation	arl.netra.ca.check.failed	Aggregate	Error
Object-store Access Denied for Aggregate Relocation During Storage Failover	gb.netra.ca.check.failed	Aggregate	Error
FabricPool Space Nearly Full	fabricpool.nearly.full	Cluster	Error
NVMe-oF Grace Period Started	nvmf.graceperiod.start	Cluster	Warning
NVMe-oF Grace Period Active	nvmf.graceperiod.active	Cluster	Warning
NVMe-oF Grace Period Expired	nvmf.graceperiod.expired	Cluster	Warning
LUN Destroyed	lun.destroy	LUN	Information
Cloud AWS MetaDataConnFail	cloud.aws.metadataConn Fail	Node	Error
Cloud AWS IAMCredsExpired	cloud.aws.iamCredsExpir ed	Node	Error
Cloud AWS IAMCredsInvalid	cloud.aws.iamCredsInvali d	Node	Error
Cloud AWS IAMCredsNotFound	cloud.aws.iamCredsNotF ound	Node	Error
Cloud AWS IAMCredsNotInitialized	cloud.aws.iamNotInitialize d	Node	Information

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
Cloud AWS IAMRoleInvalid	cloud.aws.iamRoleInvalid	Node	Error
Cloud AWS IAMRoleNotFound	cloud.aws.iamRoleNotFou nd	Node	Error
Objstore Host Unresolvable	objstore.host.unresolvable	Node	Error
Objstore InterClusterLifDown	objstore.interclusterlifDow n	Node	Error
Request Mismatch Object-store Signature	osc.signatureMismatch	Node	Error
One of NFSv4 Pools Exhausted	Nblade.nfsV4PoolExhaust	Node	Critical
QoS Monitor Memory Maxed	qos.monitor.memory.max ed	Node	Error
QoS Monitor Memory Abated	qos.monitor.memory.abat ed	Node	Information
NVMeNS Destroy	NVMeNS.destroy	Namespace	Information
NVMeNS Online	NVMeNS.offline	Namespace	Information
NVMeNS Offline	NVMeNS.online	Namespace	Information
NVMeNS Out of Space	NVMeNS.out.of.space	Namespace	Warning
Synchronous Replication Out Of Sync	sms.status.out.of.sync	SnapMirror relationship	Warning
Synchronous Replication Restored	sms.status.in.sync	SnapMirror relationship	Information
Synchronous Replication Auto Resync Failed	sms.resync.attempt.failed	SnapMirror relationship	Error
Many CIFS Connections	Nblade.cifsManyAuths	SVM	Error

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
Max CIFS Connection Exceeded	Nblade.cifsMaxOpenSam eFile	SVM	Error
Max Number of CIFS Connection Per User Exceeded	Nblade.cifsMaxSessPerU srConn	SVM	Error
CIFS NetBIOS Name Conflict	Nblade.cifsNbNameConfli ct	SVM	Error
Attempts to Connect Nonexistent CIFS Share	Nblade.cifsNoPrivShare	SVM	Critical
CIFS Shadow Copy Operation Failed	cifs.shadowcopy.failure	SVM	Error
Virus Found By AV Server	Nblade.vscanVirusDetect ed	SVM	Error
No AV Server Connection for Virus Scan	Nblade.vscanNoScanner Conn	SVM	Critical
No AV Server Registered	Nblade.vscanNoRegdSca nner	SVM	Error
No Responsive AV Server Connection	Nblade.vscanConnInactiv e	SVM	Information
AV Server too Busy to Accept New Scan Request	Nblade.vscanConnBackPr essure	SVM	Error
Unauthorized User Attempt to AV Server	Nblade.vscanBadUserPriv Access	SVM	Error
FlexGroup Constituents Have Space Issues	flexgroup.constituents.hav e.space.issues	Volume	Error
FlexGroup Constituents Space Status All OK	flexgroup.constituents.spa ce.status.all.ok	Volume	Information
FlexGroup Constituents Have Inodes Issues	flexgroup.constituents.hav e.inodes.issues	Volume	Error

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
FlexGroup Constituents Inodes Status All OK	flexgroup.constituents.ino des.status.all.ok	Volume	Information
Volume Logical Space Nearly Full	monitor.vol.nearFull	Volume	Warning
Volume Logical Space Full	monitor.vol.full	Volume	Error
Volume Logical Space Normal	monitor.vol.one.ok	Volume	Information
WAFL Volume AutoSize Fail	wafl.vol.autoSize.fail	Volume	Error
WAFL Volume AutoSize Done	wafl.vol.autoSize.done	Volume	Information

Subscribing to ONTAP EMS events

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

Before you begin

Do not subscribe to EMS events that are already added to Unified Manager automatically as this can cause confusion when receiving two events for the same issue.

About this task

You can subscribe to any number of EMS events. All the events to which you subscribe are validated, and only the validated events are applied to the clusters you are monitoring in Unified Manager. The *ONTAP 9 EMS Event Catalog* provides detailed information for all of the EMS messages for the specified version of ONTAP 9 software. Locate the appropriate version of the *EMS Event Catalog* from the ONTAP 9 Product Documentation page for a list of the applicable events.

ONTAP 9 Product Library

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



If you do not receive the ONTAP EMS events to which you have subscribed, there might be an issue with the DNS configuration of the cluster which is preventing the cluster from reaching the Unified Manager server. To resolve this issue, the cluster administrator must correct the DNS configuration of the cluster, and then restart Unified Manager. Doing so will flush the pending EMS events to the Unified Manager server.

Steps

- 1. In the left navigation pane, click Configuration > Manage Events.
- In the Configuration/Manage Events page, click the Subscribe to EMS events button.
- In the Subscribe to EMS events dialog box, enter the name of the ONTAP EMS event to which you want to subscribe.

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

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

4. Click Add.

The EMS event is added to the Subscribed EMS events list, but the Applicable to Cluster column displays the status as "Unknown" for the EMS event that you added.

- 5. Click Save and Close to register the EMS event subscription with the cluster.
- Click Subscribe to EMS events again.

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

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

After you finish

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

Managing SAML authentication settings

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

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

Identity provider requirements

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

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

Supported encryption standards

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

Validated identity providers

- Shibboleth
- Active Directory Federation Services (ADFS)

ADFS configuration requirements

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

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

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

http://www.torivar.com/2016/03/22/adfs-3-0-disable-revocation-check-windows-2012-r2/

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

Other configuration requirements

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

Enabling SAML authentication

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

Before you begin

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

About this task

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

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



Unified Manager is restarted automatically after you complete the SAML configuration on this page.

Steps

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

The fields required to configure the IdP connection are displayed.

4. Enter the IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP server.

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

5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.

You can configure the IdP server with this information at this time.

6. Click Save.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

7. Click Confirm and Logout and Unified Manager is restarted.

Results

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

After you finish

If not already completed, access your IdP and enter the Unified Manager server URI and metadata to complete the configuration.



When using ADFS as your identity provider, the Unified Manager GUI does not honor the ADFS timeout and will continue to work until the Unified Manager session timeout is reached. When Unified Manager is deployed on Windows, Red Hat, or CentOS, you can change the GUI session timeout using the following Unified Manager CLI command: um option set absolute.session.timeout=00:15:00This command sets the Unified Manager GUI session timeout to 15 minutes.

Configuring database backup settings

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

Before you begin

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

It is recommended that you use a remote location that is external to the Unified Manager host system.

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

About this task

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

Steps

- 1. In the toolbar, click , and then click Management > Database Backup.
- 2. In the Management/Database Backup page, click Actions > Database Backup Settings.
- 3. Configure the appropriate values for a backup path and retention count.

The default value for retention count is 10; you can use 0 for creating unlimited backups.

- 4. In the **Schedule Frequency** section, select the **Enable** checkbox, and then specify a daily or weekly schedule.
 - Daily

If you select this option, you must enter a time in 24-hour format for creating the backup. For example, if you specify 18:30, then a backup is created daily at 6:30 PM.

Weekly

If you select this option, you must specify the time and day for creating the backup. For example, if you specify the day as Monday and time as 16:30, then a weekly backup is created every Monday at 4:30 PM.

Click Save and Close.

Changing the local user password

You can change your local user 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 remote users cannot be changed using these steps. To change a remote user password, contact your password administrator. To change the maintenance user password, see Using the maintenance console.

Steps

- 1. Log in to Unified Manager.
- 2. From the top menu bar, click the user icon and then click **Change Password**.

The **Change Password** option is not displayed if you are a remote user.

- In the Change Password dialog box, enter the current password and the new password.
- 4. Click Save.

After you finish

If Unified Manager is configured in a high-availability configuration, you must change the password on the second node of the setup. Both instances must have same password.

Changing the Unified Manager host name

At some point, you might want to change the host name of the system on which you have installed Unified Manager. For example, you might want to rename the host to more easily identify your Unified Manager servers by type, workgroup, or monitored cluster group.

The steps required to change the host name are different depending on whether Unified Manager is running on a VMware ESXi server, on a Red Hat or CentOS Linux server, or on a Microsoft Windows server.

Changing the Unified Manager virtual appliance host name

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

Before you begin

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

About this task

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

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

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

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

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

Steps

1. Generate an HTTPS security certificate

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

2. Restart the Unified Manager virtual machine

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

Generating an HTTPS security certificate

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

Before you begin

You must have the OnCommand Administrator role.

About this task

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

Steps

- 1. In the toolbar, click , and then click HTTPS Certificate from the Setup menu.
- 2. Click Regenerate HTTPS Certificate.

The Regenerate HTTPS Certificate dialog box is displayed.

3. Select one of the following options depending on how you want to generate the certificate:

If you want to	Do this
Regenerate the certificate with the current values	Click the Regenerate Using Current Certificate Attributes option.

If you want to	Do this	
Generate the certificate using different values		the *Update the Current ficate Attributes* option.
	will use the do not en require var example,	amon Name and Alternative Names fields ne values from the existing certificate if you atter new values. The other fields do not alues, but you can enter values, for for the City, State, and Country if you want ues to be populated in the certificate.
	i	You can select the "Exclude local identifying information (e.g. localhost)" checkbox if you want to remove the local identifying information from the Alternative Names field in the certificate. When this checkbox is selected only what you enter in the field is used in the Alternative Names field. When left blank the resulting certificate will not have an Alternative Names field at all.
	+	

- 4. Click **Yes** to regenerate the certificate.
- 5. Restart the Unified Manager server so that the new certificate takes effect.

After you finish

Verify the new certificate information by viewing the HTTPS certificate.

Restarting the Unified Manager virtual machine

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

Before you begin

The virtual appliance is powered on.

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

About this task

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

Steps

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

Changing the Unified Manager host name on Linux systems

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

Before you begin

You must have root user access to the Linux system on which Unified Manager is installed.

About this task

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

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

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

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

Steps

- 1. Log in as the root user to the Unified Manager system that you want to modify.
- 2. Stop the Unified Manager software and the associated MySQL software by entering the following commands in the order shown: service ocieau stopservice ocie stopservice mysqld stop
- Change the host name using the Linux hostnamectl command: hostnamectl set-hostname new_FQDN

hostnamectl set-hostname nuhost.corp.widget.com

- 4. Regenerate the HTTPS certificate for the server:/opt/netapp/essentials/bin/cert.sh create
- 5. Restart the network service: service network restart
- 6. After the service is restarted, verify whether the new host name is able to ping itself: ping new hostname

ping nuhost

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

_	
7.	After you complete and verify your host name change, restart Unified Manager by entering the following commands in the order shown: service mysqld startservice ocie startservice ocieau start

Online Help

Introduction to OnCommand Unified Manager

OnCommand Unified Manager enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface.

Unified Manager provides the following features:

- · Discovery, monitoring, and notifications for systems that are installed with ONTAP software.
- · Dashboards to show capacity, availability, protection, and performance health of the environment.
- Enhanced alerts, events, and threshold infrastructure.
- Displays detailed graphs that plot workload activity over time; including IOPS (operations), MBps (throughput), latency (response time), utilization, performance capacity, and cache ratio.
- Identifies workloads that are overusing cluster components and the workloads whose performance is impacted by the increased activity.
- Provides suggested corrective actions that can be performed to address certain incidents and events.
- Integrates with OnCommand Workflow Automation to execute automated workflows.

Introduction to OnCommand Unified Manager health monitoring

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

The key capabilities of Unified Manager include monitoring, alerting, managing availability and capacity of clusters, managing protection capabilities, monitoring performance, configuring and managing of Infinite Volumes, annotating storage objects, and bundling of diagnostic data and sending it to technical support.

You can use Unified Manager to monitor your clusters. When issues occur in the cluster, Unified Manager notifies you about the details of such issues through events. Some events also provide you with a remedial action that you can take to rectify the issues. You can configure alerts for events so that when issues occur, you are notified through email, and SNMP traps.

You can use Unified Manager to manage storage objects in your environment by associating them with annotations. You can create custom annotations and dynamically associate clusters, storage virtual machines (SVMs), and volumes with the annotations through rules.

You can also plan the storage requirements of your cluster objects using the information provided in the capacity and health charts, for the respective cluster object.

Unified Manager health monitoring features

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

Unified Manager includes the following features:

- Discovery, monitoring, and notifications for systems that are installed with ONTAP software:
 - Physical objects: nodes, disks, disk shelves, SFO pairs, ports, and Flash Cache
 - Logical objects: clusters, storage virtual machines (SVMs), aggregates, volumes, LUNs, namespaces, qtrees, LIFs, Snapshot copies, junction paths, NFS exports, CIFS shares, user and group quotas, and initiator groups
 - Protocols: CIFS, NFS, FC, iSCSI, NVMe, and FCoE
 - Storage efficiency: SSD aggregates, Flash Pool aggregates, FabricPool aggregates, deduplication, and compression
 - Protection: SnapMirror relationships (synchronous and asynchronous) and SnapVault relationships
- Viewing the cluster discovery and monitoring status
- MetroCluster configuration: viewing and monitoring the configuration, MetroCluster switches and bridges, issues, and connectivity status of the cluster components
- · Enhanced alerts, events, and threshold infrastructure
- · LDAP, LDAPS, SAML authentication, and local user support
- RBAC (for a predefined set of roles)
- · AutoSupport and support bundle
- · Enhanced dashboard to show capacity, availability, protection, and performance health of the environment
- Volume move interoperability, volume move history, and junction path change history
- Scope of Impact area that graphically displays the resources that are impacted for events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events
- Possible Effect area that displays the effect of the MetroCluster events
- Suggested Corrective Actions area that displays the actions that can be performed to address events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events
- Resources that Might be Impacted area that displays the resources that might be impacted for events such as for the Volume Offline event, the Volume Restricted event, and the Thin-Provisioned Volume Space At Risk event
- · Support for SVMs with:
 - FlexVol volumes
 - · FlexGroup volumes
 - Infinite Volumes
- Support for monitoring node root volumes
- Enhanced Snapshot copy monitoring, including computing reclaimable space and deleting Snapshot copies
- Annotations for storage objects
- Report creation and management of storage object information such as physical and logical capacity, utilization, space savings, and related events
- Integration with OnCommand Workflow Automation to execute workflows

The Storage Automation Store contains NetApp-certified automated storage workflow packs developed for use with OnCommand Workflow Automation (WFA). You can download the packs, and then import them to

Introduction to OnCommand Unified Manager performance monitoring

OnCommand Unified Manager provides performance monitoring capabilities and event root-cause analysis for systems that are running NetApp ONTAP software.

Unified Manager helps you to identify workloads that are overusing cluster components and decreasing the performance of other workloads on the cluster. By defining performance threshold policies you can also specify maximum values for certain performance counters so that events are generated when the threshold is breached. Unified Manager alerts you about these performance events so that you can take corrective action, and bring performance back to normal levels of operation. You can view and analyze events in the Unified Manager UI.

Unified Manager monitors the performance of two types of workloads:

· User-defined workloads

These workloads consist of FlexVol volumes and FlexGroup volumes that you have created in your cluster.

· System-defined workloads

These workloads consist of internal system activity.

Unified Manager performance monitoring features

Unified Manager collects and analyzes performance statistics from systems running ONTAP software. It uses dynamic performance thresholds and user-defined performance thresholds to monitor a variety of performance counters over many cluster components.

A high response time (latency) indicates that the storage object, for example, a volume, is performing slower than normal. This issue also indicates that the performance has decreased for client applications that are using the volume. Unified Manager identifies the storage component where the performance issue lies and provides a list of suggested actions you can take to address the performance issue.

Unified Manager includes the following features:

- Monitors and analyzes workload performance statistics from a system running ONTAP software.
- Tracks performance counters for clusters, nodes, aggregates, ports, SVMs, volumes, LUNs, NVMe namespaces, and LIFs.
- Displays detailed graphs that plot workload activity over time; including IOPS (operations), MBps (throughput), latency (response time), utilization, performance capacity, and cache ratio.
- Enables you to create user-defined performance threshold policies that trigger events and send email alerts when the thresholds are breached.
- Uses system-defined thresholds and dynamic performance thresholds that learn about your workload activity to identify and alert you to performance issues.
- Clearly identifies the cluster component that is in contention.
- Identifies workloads that are overusing cluster components and the workloads whose performance is impacted by the increased activity.

What the Unified Manager server does

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

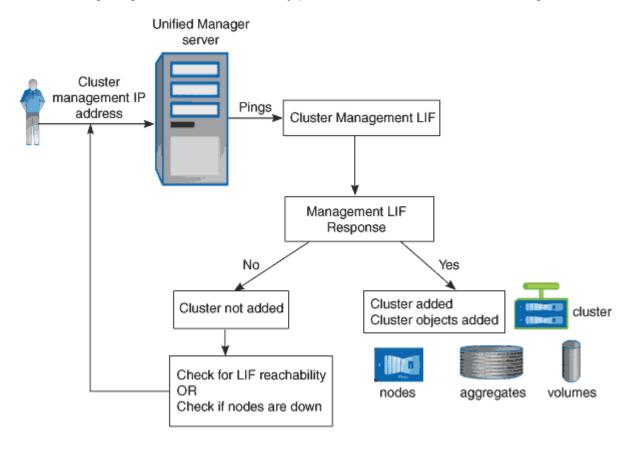
Unified Manager collects cluster information, stores the data in the database, and analyzes the data to see if there are any cluster issues.

How the discovery process works

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

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

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



Cluster configuration and performance data collection activity

The collection interval for *cluster configuration data* is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster details in the Unified Manager UI. This interval applies when making changes to a cluster too.

For example, if you add two new volumes to an SVM in a cluster, you see those new objects in the UI after the next polling interval, which could be up to 15 minutes.

Unified Manager collects current *performance statistics* from all monitored clusters every five minutes. It analyzes this data to identify performance events and potential issues. It retains 30 days of five-minute historical performance data and 390 days of one-hour historical performance data. This enables you to view very granular performance details for the current month, and general performance trends for up to a year.

The collection polls are offset by a few minutes so that data from every cluster is not sent at the same time, which could affect performance.

The following table describes the collection activities that Unified Manager performs:

Activity	Time interval	Description
Performance statistics poll	Every 5 minutes	Collects real-time performance data from each cluster.
Statistical analysis	Every 5 minutes	After every statistics poll, Unified Manager compares the collected data against user-defined, system-defined, and dynamic thresholds. If any performance thresholds have been breached, Unified Manager generates events and sends email to specified users, if configured to do so.
Configuration poll	Every 15 minutes	Collects detailed inventory information from each cluster to identify all the storage objects (nodes, SVMs, volumes, and so on).
Summarization	Every hour	Summarizes the latest 12 five- minute performance data collections into hourly averages. The hourly average values are used in some of the UI pages, and they are retained for 390 days.
Forecast analysis and data pruning	Every day after midnight	Analyzes cluster data to establish dynamic thresholds for volume latency and IOPS for the next 24 hours. Deletes from the database any five-minute performance data older than 30 days.
Data pruning	Every day after 2 a.m.	Deletes from the database any events and dynamic thresholds older than 390 days.

Activity	Time interval	Description
Data pruning	Every day after 3:30 a.m.	Deletes from the database any one-hour performance data older than 390 days.

What a data continuity collection cycle is

A data continuity collection cycle retrieves performance data outside of the real-time cluster performance collection cycle that runs, by default, every five minutes. Data continuity collections enable Unified Manager to fill in gaps of statistical data that occur when it was unable to collect real-time data.

Data continuity collection is supported only on clusters installed with ONTAP version 8.3.1 or later software.

Unified Manager performs data continuity collection polls of historical performance data when the following events occur:

A cluster is initially added to Unified Manager.

Unified Manager gathers historical performance data for the previous 15 days. This enables you to view two weeks of historical performance information for a cluster a few hours after it is added.

Additionally, system-defined threshold events are reported for the previous period, if any exist.



15 days of historical volume statistics are not currently collected.

• The current performance data collection cycle does not finish on time.

If the real-time performance poll goes beyond the five-minute collection period, a data continuity collection cycle is initiated to gather that missing information. Without the data continuity collection, the next collection period is skipped.

- Unified Manager has been inaccessible for a period of time and then it comes back online, as in the following situations:
 - It was restarted.
 - It was shut down during a software upgrade or when creating a backup file.
 - A network outage is repaired.
- A cluster has been inaccessible for a period of time and then it comes back online, as in the following situations:
 - A network outage is repaired.
 - · A slow wide area network connection delayed the normal collection of performance data.

A data continuity collection cycle can collect a maximum of 24 hours of historical data. If Unified Manager is down for longer than 24 hours, a gap in performance data appears in the UI pages.

A data continuity collection cycle and a real-time data collection cycle cannot run at the same time. The data continuity collection cycle must finish before the real-time performance data collection is initiated. When the data continuity collection is required to collect more than one hour of historical data, then you see a banner message for that cluster at the top of the Performance dashboard.

What the timestamp means in collected data and events

The timestamp that appears in collected health and performance data, or that appears as the detection time for an event, is based on the ONTAP cluster time, adjusted to the time zone set on the web browser.

It is highly recommended that you use a Network Time Protocol (NTP) server to synchronize the time on your Unified Manager servers, ONTAP clusters, and web browsers.



If you see timestamps that look incorrect for a particular cluster, you might want to check that the cluster time has been set correctly.

Understanding the user interface

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

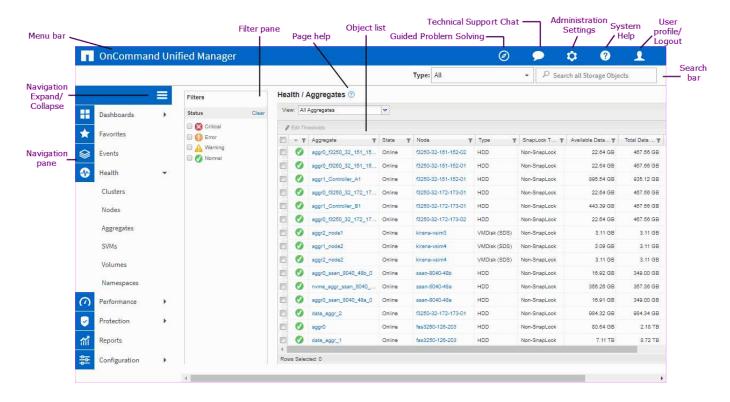
You can select a preferred view and use the action buttons as necessary. Your screen configuration is saved in a workspace so that all of the functionality you require is available when you start Unified Manager. However, when you navigate from one view to another, and then navigate back, the view might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

Typical window layouts

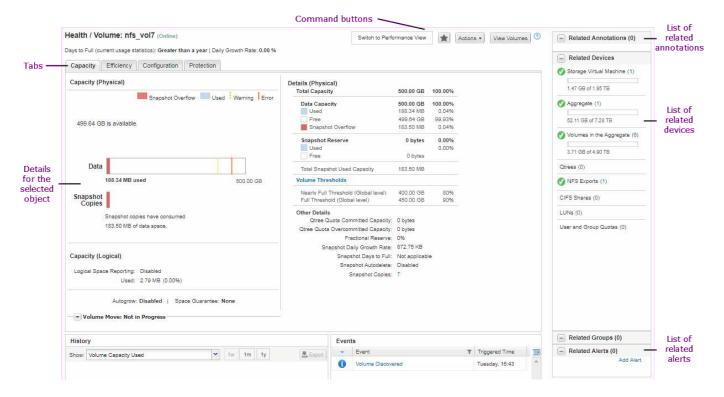
Understanding the typical window layouts helps you to navigate and use OnCommand Unified Manager effectively. Most Unified Manager windows are similar to one of two general layouts: object list or details. The recommended display setting is at least 1280 by 1024 pixels.

Not every window contains every element in the following diagrams.

Object list window layout



Object details window layout



Window layout customization

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

Sorting

You can click the column header to change the sort order of the column entries. When you click the column header, the sort arrows (* and *) appear for that column.

Filtering

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

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

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

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

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

· Hiding or redisplaying the columns

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

Exporting data

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

Using the Unified Manager Help

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

About this task

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

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

Choices

- To learn about specific fields or parameters, click ...
- To view all the Help contents, click (> Help/Documentation in the menu bar.

You can find more detailed information by expanding any portion of the Table of Contents in the navigation pane.

- To search the Help contents, click the **Search** tab in the navigation pane, type the word or series of words you want to find, and click **Go!**
- To print Help topics, click the printer icon.

Bookmarking your favorite Help topics

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

Steps

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

Exporting data to CSV files for reporting

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

Before you begin

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

About this task

You can export data to a .csv file from the health and performance inventory pages and from the event inventory page.

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

Steps

1. Perform one of the following actions:

If you want to export	Do this
Event details	Click Events.

If you want to export	Do this
Storage object inventory details	Click Health or Performance from the left- navigation menu, and then select a storage object.
Storage capacity and protection history details	Click Health > Aggregates or Health > Volumes , then select a single aggregate or volume.
Qtree or NFS Exports information for an SVM	Click Health > SVMs , select a single SVM, and then select the Qtrees or NFS Exports tab.
Storage object top 10 performance details	Click Performance > Clusters , then select a cluster and choose the Top Performers tab. Then select a storage object and performance counter.

- 2. Click the **Export** button.
- 3. Click Export to CSV to confirm the export request.

From the Top Performers tab, and from the SVMs details page, you can choose to download a report of the statistics for the single cluster you are viewing or for all clusters in the data center.

The file is downloaded.

4. Open the .csv file in the appropriate application.

Searching for storage objects

To quickly access a specific object, you can use the **Search all Storage Objects** field at the top-right of the interface. This method of global search across all objects enables you to quickly locate specific objects by type. Search results are sorted by storage object type and you can filter them further by object using the **Type** drop-down menu.

Before you begin

- You must have one of the following roles to perform this task: Operator, OnCommand Administrator, or Storage Administrator.
- A valid search must contain at least three characters.

About this task

When using the Type drop-down menu value "All", the global search displays the total number of results found in all object categories; with a maximum of 25 search results for each object category. You can select a specific object type from the Type drop-down menu to refine the search within a specific object type. In this case the returned list is not restricted to the top 25 objects.

The object types you can search for include:

- Clusters
- Nodes

- SVMs
- Aggregates
- Volumes
- Qtrees
- CIFS Shares
- · User or Group Quotas
- LUNs
- NVMe Namespaces
- Initiator Groups
- Initiators

You can click any object in the search results to navigate to the Health details page for that object. If there is no direct health page for an object, then the Health page of the parent object is displayed. For example, when searching for a specific LUN, the SVM details page on which the LUN resides is displayed.

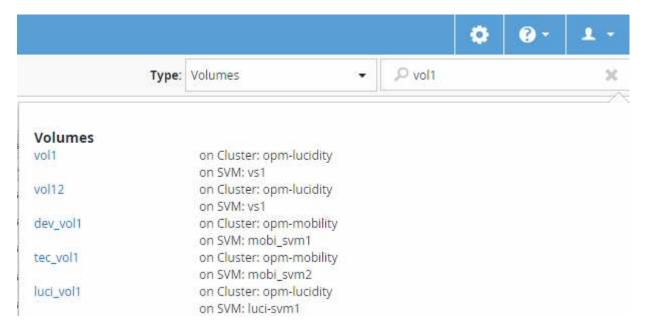


Ports are not searchable in the global search bar.

Steps

- 1. Select an object type from the **Type** menu to refine the search results for only a single object type.
- 2. Type a minimum of three characters of the object name in the Search all Storage Objects field.

In this example, the **Type** drop-down box has the Volumes object type selected. Typing "vol1" into the **Search all Storage Objects** field displays a list of all volumes whose names contain these characters.



Filtering performance inventory page content

You can filter performance inventory data in Unified Manager to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Unified Manager

pages to show only the results in which you are interested. This provides a very efficient method of displaying only the performance data in which you are interested.

About this task

Use **Filtering** to customize the grid view based on your preferences. Available filter options are based on the object type being viewed in the grid. If filters are currently applied, an asterisk (*) displays at the left of the Filtering control.

Four types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are contains and starts with .
Number	The operators are greater than and less than .
Resource	The operators are name contains and name starts with.
Status	The operators are is and is not .

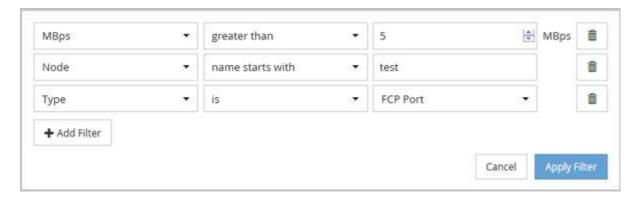
All three fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

You can add filters using the Filtering panel.

- 1. At the top of the page, click **Filtering**. The Filtering panel displays.
- 2. In the Filtering panel, click the left drop-down list, and select an object name: for example, *Cluster*, or a performance counter.
- Click the center drop-down list, and select the boolean operator name contains or name starts with if the
 first selection was an object name. If the first selection was a performance counter, select greater than or
 less than. If the first selection was Status, select is or is not.
- 4. If your search criteria requires a numeric value, up and down arrow buttons display in the field at the right. You can click the up and down arrow buttons to display your desired numeric value.
- 5. If required, type your non-numeric search criteria in the text field at the right.
- To add filters, click Add Filter. An additional filter field displays. Complete this filter using the process
 described in the preceding steps. Note that upon adding your fourth filter, the Add Filter button no longer
 displays.
- 7. Click **Apply Filter**. The filter options are applied to the grid and an asterisk (*) is displayed in the Filtering button.
- 8. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
- 9. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example

The illustration shows the Filtering panel with three filters. The **Add Filter** button displays when you have fewer than the maximum of four filters.



After clicking Apply Filter, the Filtering panel closes and applies your filters.

Accessing OnCommand System Manager from the Unified Manager interface

When troubleshooting requires that you make configuration changes to a cluster, you can use the System Manager graphical interface instead of the ONTAP command-line interface. System Manager is included with ONTAP as a web service, it is enabled by default, and it is accessible by using a browser.

Before you begin

You must have a cluster user account configured with the admin role and the http, ontapi, and console application types.

Steps

- 1. In the left navigation pane, click **Dashboards** > **Cluster View**.
- 2. In the Dashboards/Cluster View page, select the cluster that you want to manage.

An overview of the monitoring status, capacity, and performance for that cluster is displayed.

3. Click the System Manager icon.

If the cluster uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate on the cluster for server authentication.

4. Log in to System Manager by using your cluster administrator credentials.

If login to the System Manager user interface is protected using SAML authentication you will enter your credentials in the identity provider (IdP) login page instead of the System Manager login page.

Monitoring cluster health and performance from the dashboards

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

Three dashboards provide unique views into the health and performance of your clusters:

- The Dashboards/Overview page provides information about the health and performance of your storage objects.
- The Dashboards/Performance page provides high-level performance status of all the clusters that are being monitored.
- The Dashboards/Cluster View page provides information about individual clusters.

Understanding the Health overview dashboard

The Unified Manager Health overview dashboard provides cumulative at-a-glance information about the health of your storage and virtualized environment. The Dashboards/Overview page provides health information about your storage objects separated into four health categories; availability, capacity, performance, and protection of the storage objects.

The following image illustrates the panes that are displayed on Dashboards/Overview page:









Health Overview area

Displays, as a graph, information about the health of your storage objects such as clusters, aggregates, and storage virtual machines (SVMs), and the health of your protection relationships. The Dashboards/Overview page displays events generated for the following categories:

Clicking on the number above any of the yellow or red bar charts displays the Events inventory page including only those events. Clicking on the number below any of the charts displays the object inventory page including only those objects.

Availability

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

Capacity

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

Performance

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

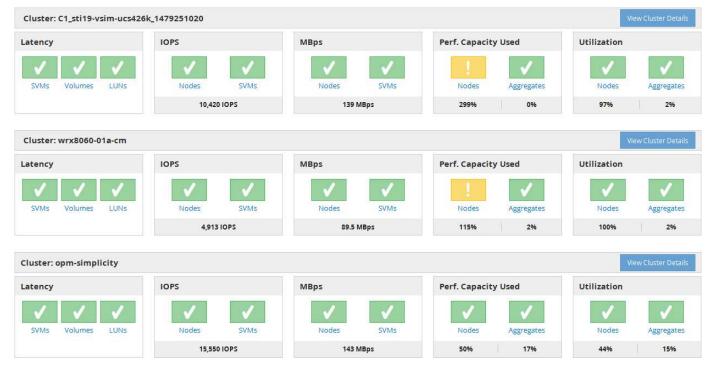
Protection

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

Understanding the Performance dashboard

The Unified Manager Performance dashboard provides a high-level overview of the performance status for all the clusters that are being monitored in your environment. Clusters that have performance issues are ordered at the top of the page by severity. The information on the dashboard is updated automatically at each five-minute performance collection period.

The following image shows an example of a Unified Manager Performance dashboard that is monitoring two clusters:



The status icons that represent the storage objects can be in the following states, sorted from highest severity to lowest severity:

- Critical (X): One or more new critical performance events have been reported for the object.
- Warning (¹/₁): One or more new warning performance events have been reported for the object.
- Normal (
): No new performance events have been reported for the object.



The color indicates whether new events exist for the object. Events that are no longer active, called obsolete events, do not affect the color of the icon.

Cluster performance counters

The following performance categories are displayed for each cluster:

Latency

Shows how quickly the cluster is responding to client application requests, in milliseconds per operation.

IOPS

Shows the operating speed of the cluster, in number of input/output operations per second.

• MBps

Shows how much data is being transferred to and from the cluster, in megabytes per second.

· Performance Capacity Used

Shows whether any nodes or aggregates are overusing their available performance capacity.

Utilization

Shows whether the resources on any nodes or aggregates are being overused.

To analyze the performance of your cluster and storage objects, you can perform one of the following actions:

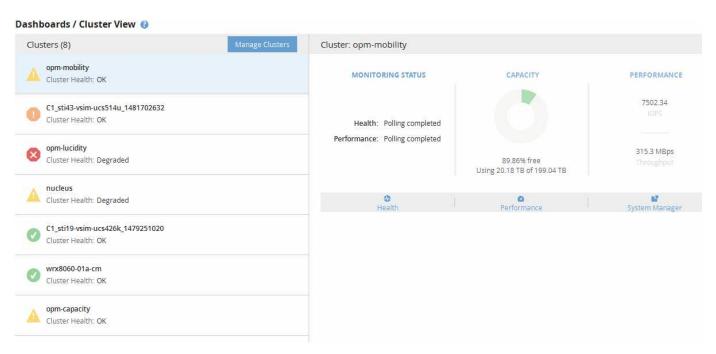
- You can click **View Cluster Details** to display the Cluster Landing page, where you can view detailed performance and event information for the selected cluster and storage objects.
- You can click one of the red or yellow status icons of an object to display the Inventory page for that object, where you can view details about the storage object.

For example, clicking a volume icon displays the Performance/Volume inventory page with a list of all the volumes in the selected cluster, sorted from worst performance to best performance.

Understanding the Cluster View dashboard

The Unified ManagerCluster View overview dashboard provides high-level information about the health of the clusters you are managing. The Cluster View dashboard consists of two major sections: Managed Clusters (on the left) and Cluster Details (on the right).

The following image shows an example of a Unified ManagerCluster View dashboard that is monitoring eight clusters:



The status icon next to each cluster name can be in the following states:

- Critical (🔞): One or more active critical events have been reported for the cluster.
- Error (1): One or more active error events have been reported for the cluster.
- Warning (1): One or more active warning events have been reported for the cluster.
- Normal (): No active events have been reported for the cluster.



The color indicates whether active (new or acknowledged) events exist for the object. Events that are no longer active, called obsolete events, do not affect the color of the icon.

To display additional information about a cluster, you can perform one of the following actions:

- You can click a cluster name to display overview information on the monitoring status, capacity status, and performance status of the cluster.
- You can click **Manage Clusters** to display the Configuration/Cluster Data Sources page, where you can view detailed status information for all the clusters being managed by this instance of Unified Manager.

Description of dashboard windows

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

Dashboards/Overview page

The Dashboards/Overview page displays, as a graph, the health of storage objects such as clusters, aggregates, and storage virtual machines (SVMs). Based on the availability, capacity, performance, and protection-related events that are generated, these storage objects are categorized as Healthy, At Risk, or Have Incidents, or as Healthy, Warning or Error for protection-related events.

Cluster not reachable bar

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

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

You can view detailed information about a cluster that is unreachable by clicking the **Details** button. This action opens the Events inventory page. After the bar is closed, it is displayed again only when you log back in to Unified Manager.

Overview area

Availability pane

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

This pane also displays the number of storage objects in each of the categories. Clicking on any of the object totals takes you to the page for that object. For example, clicking the cluster total takes you to the Health/Clusters inventory page. Clicking on the numbers at the top of a column takes you to the Events inventory page.

Capacity pane

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

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

Performance pane

Displays information about the performance of clusters, SVMs, and volumes that are monitored by Unified Manager. Based on the performance-related incidents that are generated, the storage objects are categorized as Healthy, At Risk, or Have Incidents. For example, the status of a volume whose I/O response time to its workload has reached the maximum threshold value is displayed as Have Incidents.

This pane also displays the total number of clusters, SVMs, and volumes that are monitored by Unified Manager. Clicking on any of the object totals takes you to the page for that object. For example, clicking the cluster total takes you to the Performance/Cluster inventory page.

Protection pane

Displays information about protection relationships that are monitored by Unified Manager. The protection relationships are categorized as Healthy, Warning, or Error. For example, a relationship that has a lag duration that exceeds the lag warning threshold is displayed as Warning.

This pane also displays the total number of storage objects in each of the protection categories. Clicking the links for the Lag Status, Asynchronous Vault, Asynchronous Mirror, or Synchronous categories takes you to a filtered list of those objects in the Protection/Volume Relationships page.

Dashboards/Performance page

You can use the Unified Manager Performance Dashboard to view the high-level performance status of all the clusters that are being monitored. The Dashboards/Performance page also displays a banner message when Unified Manager is unable to communicate with a cluster that it is monitoring.

Overview

The clusters are ordered based on severity using the following criteria:

- 1. If a cluster is unreachable.
- 2. If a cluster has one or more active critical performance events (red object icon).
- 3. If a cluster has one or more active warning performance events (yellow object icon).
- If clusters have no active performance events (green object icon): the clusters are sorted by highest IOPS.

Critical events are generated when a critical limit in a user-defined performance threshold policy is exceeded. Warning events are sent when a warning limit in a user-defined performance threshold policy is exceeded, or when a system-defined threshold policy or dynamic threshold is exceeded.



The sort order is determined by the total number of active (new or acknowledged) events, not by the number of objects that have events. For example, if Cluster A has seven critical volume latency events, and Cluster B has two critical volume latency events and two critical node IOPS events (for a total of four critical events), Cluster A (with one red object icon) would appear higher on the list, even though Cluster B has two red object icons.

The following commonly monitored event types are displayed for each cluster:

• Latency events for storage virtual machines (SVMs), volumes, and LUNs

- · IOPS events for nodes and SVMs
- · MBps events for nodes and SVMs
- · Performance capacity used events for nodes and aggregates
- · Utilization events for nodes and aggregates

Unified Manager can receive performance events for other storage objects and counters—for example, MBps events for aggregates, and IOPS events for volumes. If a cluster has these types of events, the icon next to the cluster name indicates a warning (yellow) or critical (red) event. This icon might indicate that events exist when none of the five counter panel icons are yellow or red because these event types do not fit into the existing counter panel categories.

You can click the object icon to display the Performance Inventory page for that object, where you can view all objects of that type in this cluster. For example, clicking a volume icon displays the Performance/Volume Inventory page, showing a list of all the volumes in the selected cluster, sorted from worst to best performance.

Performance counters

There are five performance counters, each displayed in a separate panel, for every cluster. This information is updated automatically after each five-minute collection period:

Counter	Description
Latency	Shows how quickly the cluster is responding to client application requests, in milliseconds per operation (ms/op). The icon area indicates whether any SVMs, volumes, or LUNs have any active events based on the latency value crossing a threshold setting.
IOPS	 Shows the operating speed of the storage system, in number of input/output operations per second (IOPS). The icon area indicates whether any nodes or SVMs have any active events based on the number of IOPS crossing a threshold setting. The bottom area displays the total cluster IOPS for the last five-minute collection period.
MBps	 Shows how much data is being transferred to and from the cluster, in megabytes per second (MBps). The icon area indicates whether any nodes or SVMs have any active events based on the MBps value crossing a threshold setting. The bottom area displays the total cluster throughput for the last five-minute collection period.

Counter	Description
Performance Capacity Used	 Shows whether any nodes or aggregates are overusing their available performance capacity. The icon area indicates whether any nodes or aggregates have any active events based on the performance capacity used value crossing a threshold setting. The bottom area displays the highest performance capacity used value from the busiest node and busiest aggregate. Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.
Utilization	Shows whether the resources on any nodes or aggregates are being overused. • The icon area indicates whether any nodes or aggregates have any active events based on the utilization value crossing a threshold setting. • The bottom area displays the highest utilization value from the busiest node and busiest aggregate.

An ellipsis (...) in the header area indicates that performance data is currently being collected.

You can click **View Cluster Details** for a cluster that has performance events to display the Performance Cluster Landing page, where you can view detailed performance information about the cluster and other storage objects.

Cluster status messages

If a cluster that is managed by Unified Manager becomes unavailable, a status message banner is displayed above the performance counters. A **Details** button is displayed at the right of the status message banner if the cluster is unreachable. By clicking the **Details** button in the status message, you can navigate to the Cluster Data Sources page, which shows complete information about the issue. On the Cluster Data Sources page, you can find the data that is required to troubleshoot the issue that made the cluster become unavailable.

Dashboards/Cluster View page

The Dashboards/Cluster View page displays overview information about the clusters you are managing.

Clicking Manage Clusters takes you to the Configuration/Cluster Data Sources page.

Clicking on a cluster displays overview information on the monitoring status, capacity status, and performance status of the cluster.

The Cluster View dashboard consists of two major sections: Managed Clusters (on the left) and Cluster Details (on the right).

Managed Clusters section

Lists all the clusters the Unified Manager is monitoring. The following details are provided for each cluster in the list:

- Cluster status icon: The status can be Critical (♥), Error (♠), Warning (♠), or Normal (♥).
- IP address or host name: Provides the host name of the cluster, and the IP address or FQDN.
- Cluster Health: Provides information about the health of the cluster as monitored by Unified Manager.

The health status can have one of the following values: OK, OK with suppressed, Degraded, and Components not reachable.

Cluster Details section

Provides information about the monitoring status, capacity, and performance of the selected cluster.

• Monitoring Status: Displays the ongoing health and performance monitoring status.

The monitoring status can have the following values: Discovering, Poll completed, Poll failed, or Not available. The monitoring status displays an error message when the corresponding monitoring job (health or performance) fails.

- Capacity: Displays the total, used, and free storage capacity of the selected cluster.
- Performance: Displays the average operating speed of the cluster in number of IOPS (input/output operations per second), and the average throughput of the selected cluster in MBps (megabytes per second).

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

- The Health link navigates to the Health/Cluster details page of the selected cluster.
- The Performance link navigates to the Performance/Cluster details page of the selected cluster.
- The System Manager link navigates to the login page for OnCommand System Manager so that you can manage cluster settings.

Managing storage objects using the Favorites option

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

Tasks you can perform from the Favorites dashboard

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

Viewing the Favorites list

You can view the capacity, performance, and protection details of selected storage objects from the Favorites list. The details of a maximum of 20 storage objects are displayed in the Favorites list.

Adding storage objects to the Favorites list

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

Removing storage objects from the Favorites list

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

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

You can add storage objects to a Favorites list so you can monitor the objects for health, capacity, and performance. You can use object status in the Favorites list to determine issues and fix them before they become critical. The Favorites list also provides the most recent monitoring status of a storage object. You can remove storage objects from the Favorites list when you no longer require them to be marked as favorite.

About this task

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

Steps

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

Adding an aggregate to the Favorites list

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

After you finish

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

Cluster favorite card

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

Cluster attributes

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

· Cluster health status

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

Cluster name

Name of the cluster.

Capacity

Total free space on the cluster.

Configuration

Configuration details of the cluster.

IP Address

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

Number of nodes

Number of nodes in the cluster.

Performance

Performance details of the cluster.

• IOPS

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

Throughput

Average throughput over the last 72 hours, in MBps.

Aggregate favorite card

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

Aggregate attributes

The Aggregate favorite card displays the following aggregate attributes:

Aggregate health status

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

Aggregate name

Name of the aggregate.

Position your cursor over the aggregate name to view the name of the cluster to which the aggregate belongs.

Capacity

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

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

Performance

Performance details of the aggregate.

IOPS

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

Throughput

Average throughput over the last 72 hours, in MBps.

Latency

Average response time required for an operation, in milliseconds.

Volume favorite card

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

Volume attributes

The Volume favorite card displays the following volume attributes:

· Volume health status

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

Volume name

Name of the volume.

Capacity

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

Protection

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

Performance

Performance statistics for the volume.

• IOPS

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

Throughput

Average throughput over the last 72 hours, in MBps.

Latency

Average response time required for an operation, in milliseconds.

Managing events

Events help you to identify issues in the clusters that are monitored.

What health events are

Health events are notifications that are generated automatically when a predefined condition occurs or when an object crosses a health 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.

Health 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.

What performance events are

Performance events are incidents related to workload performance on a cluster. They help you identify workloads with slow response times. Together with health events that occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

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

What happens when an event is received

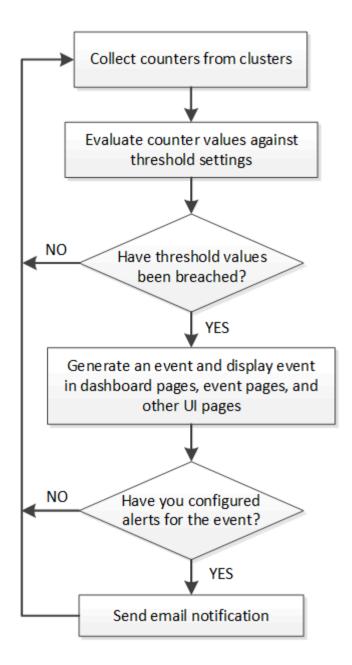
When Unified Manager receives an event, it is displayed in the Dashboards/Overview page, in the Summary and Explorer tabs of the Performance/Cluster page, in the Events inventory page, and in the object-specific inventory page (for example, the Health/Volumes inventory page).

When Unified Manager detects multiple continuous occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events. The duration of the event is incremented to indicate that the event is still active.

Depending on how you configure settings in the Configuration/Alerting page, you can notify other users about these events. The alert causes the following actions to be initiated:

- An email about the event can be sent to all Unified Manager Administrator users.
- The event can be sent to additional email recipients.
- An SNMP trap can be sent to the trap receiver.
- A custom script can be executed to perform an action.

This workflow is shown in the following diagram.



Configuration changes detected by Unified Manager

Unified Manager monitors your clusters for configuration changes to help you determine whether a change might have caused or contributed to a performance event. The Performance Explorer pages display a change event icon () to indicate the date and time when the change was detected.

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

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

A volume moves between aggregates.

Unified Manager can detect when the move is in progress, completed, or failed. If Unified Manager is down

during a volume move, when it is back up it detects the volume move and displays a change event for it.

 The throughput (MBps or IOPS) limit of a QoS policy group that contains one or more monitored workloads changes.

Changing a policy group limit can cause intermittent spikes in the latency (response time), which might also trigger events for the policy group. The latency gradually returns back to normal and any events caused by the spikes become obsolete.

• A node in an HA pair takes over or gives back the storage of its partner node.

Unified Manager can detect when the takeover, partial takeover, or giveback operation has been completed. If the takeover is caused by a panicked node, Unified Manager does not detect the event.

An ONTAP upgrade or revert operation is completed successfully.

The previous version and new version are displayed.

Configuring event retention settings

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

Before you begin

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

About this task

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

Steps

- 1. In the left navigation pane, click **Configuration > Manage Events**.
- In the Configuration/Manage Events page, click the Event Retention Settings button.
- Configure the appropriate settings in the Event Retention Settings dialog box.
- 4. Click Save and Close.

Configuring event notification settings

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

Before you begin

You must have the following information:

· Email address from which the alert notification is sent

The email address appears in the "From" field in sent alert notifications. If the email cannot be delivered for any reason, this email address is also used as the recipient for undeliverable mail.

- SMTP server host name, and the user name and password to access the server
- SNMP version, trap destination host IP address, outbound trap port, and the community to configure the SNMP trap

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

Notes:

- If the From Address is pre-filled with the address "OnCommand@localhost.com", you should change it to a real, working email address to make sure that all email notifications are delivered successfully.
- If the host name of the SMTP server cannot be resolved, you can specify the IP address (IPv4 or IPv6)
 of the SMTP server instead of the host name.

What Event Management System events are

The Event Management System (EMS) collects event data from different parts of the ONTAP kernel and provides event forwarding mechanisms. These ONTAP events can be reported as EMS events in Unified Manager. Centralized monitoring and management eases configuration of critical EMS events and alert notifications based on these EMS events.

The Unified Manager address is added as a notification destination to the cluster when you add the cluster to Unified Manager. An EMS event is reported as soon as the event occurs in the cluster.

There are two methods for receiving EMS events in Unified Manager:

- A certain number of important EMS events are reported automatically.
- You can subscribe to receive individual EMS events.

The EMS events that are generated by Unified Manager are reported differently depending on the method in which the event was generated:

Functionality	Automatic EMS messages	Subscribed EMS messages
Available EMS events	Subset of EMS events	All EMS events

Functionality	Automatic EMS messages	Subscribed EMS messages
EMS message name when triggered	Unified Manager event name (converted from EMS event name)	Non-specific in the format "Error EMS received". The detailed message provides the dot-notation format of the actual EMS event
Messages received	As soon as the cluster has been discovered	After adding each required EMS event to Unified Manager, and after the next 15 minute polling cycle
Event life cycle	Same as other Unified Manager events: New, Acknowledged, Resolved, and Obsolete states	The EMS event is made obsolete after the cluster is refreshed, after 15 minutes, from when the event was created
Captures events during Unified Manager downtime	Yes, when the system starts up it communicates with each cluster to acquire missing events	No
Event details	Suggested corrective actions are imported directly from ONTAP to provide consistent resolutions	Corrective actions not available in Event Details page



Some of the new automatic EMS events are Informational events that indicate that a previous event has been resolved. For example, the "FlexGroup Constituents Space Status All OK" Informational event indicates that the "FlexGroup Constituents Have Space Issues" Error event has been resolved. Informational events cannot be managed using the same event life cycle as other event severity types, however, the event is obsoleted automatically if the same volume receives another "Space Issues" Error event.

EMS events that are added automatically to Unified Manager

When using Unified Manager 9.4 or greater software, the following ONTAP EMS events are added automatically to Unified Manager. These events will be generated when triggered on any cluster that Unified Manager is monitoring.

The following EMS events are available when monitoring clusters running ONTAP 9.5 or greater software:

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
Object-store Access Denied for Aggregate Relocation	arl.netra.ca.check.failed	Aggregate	Error

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
Object-store Access Denied for Aggregate Relocation During Storage Failover	gb.netra.ca.check.failed	Aggregate	Error
FabricPool Space Nearly Full	fabricpool.nearly.full	Cluster	Error
NVMe-oF Grace Period Started	nvmf.graceperiod.start	Cluster	Warning
NVMe-oF Grace Period Active	nvmf.graceperiod.active	Cluster	Warning
NVMe-oF Grace Period Expired	nvmf.graceperiod.expired	Cluster	Warning
LUN Destroyed	lun.destroy	LUN	Information
Cloud AWS MetaDataConnFail	cloud.aws.metadataConn Fail	Node	Error
Cloud AWS IAMCredsExpired	cloud.aws.iamCredsExpir ed	Node	Error
Cloud AWS IAMCredsInvalid	cloud.aws.iamCredsInvali d	Node	Error
Cloud AWS IAMCredsNotFound	cloud.aws.iamCredsNotF ound	Node	Error
Cloud AWS IAMCredsNotInitialized	cloud.aws.iamNotInitialize d	Node	Information
Cloud AWS IAMRoleInvalid	cloud.aws.iamRoleInvalid	Node	Error
Cloud AWS IAMRoleNotFound	cloud.aws.iamRoleNotFou nd	Node	Error
Objstore Host Unresolvable	objstore.host.unresolvable	Node	Error
Objstore InterClusterLifDown	objstore.interclusterlifDow n	Node	Error

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
Request Mismatch Object-store Signature	osc.signatureMismatch	Node	Error
One of NFSv4 Pools Exhausted	Nblade.nfsV4PoolExhaust	Node	Critical
QoS Monitor Memory Maxed	qos.monitor.memory.max ed	Node	Error
QoS Monitor Memory Abated	qos.monitor.memory.abat ed	Node	Information
NVMeNS Destroy	NVMeNS.destroy	Namespace	Information
NVMeNS Online	NVMeNS.offline	Namespace	Information
NVMeNS Offline	NVMeNS.online	Namespace	Information
NVMeNS Out of Space	NVMeNS.out.of.space	Namespace	Warning
Synchronous Replication Out Of Sync	sms.status.out.of.sync	SnapMirror relationship	Warning
Synchronous Replication Restored	sms.status.in.sync	SnapMirror relationship	Information
Synchronous Replication Auto Resync Failed	sms.resync.attempt.failed	SnapMirror relationship	Error
Many CIFS Connections	Nblade.cifsManyAuths	SVM	Error
Max CIFS Connection Exceeded	Nblade.cifsMaxOpenSam eFile	SVM	Error
Max Number of CIFS Connection Per User Exceeded	Nblade.cifsMaxSessPerU srConn	SVM	Error
CIFS NetBIOS Name Conflict	Nblade.cifsNbNameConfli ct	SVM	Error
Attempts to Connect Nonexistent CIFS Share	Nblade.cifsNoPrivShare	SVM	Critical

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
CIFS Shadow Copy Operation Failed	cifs.shadowcopy.failure	SVM	Error
Virus Found By AV Server	Nblade.vscanVirusDetect ed	SVM	Error
No AV Server Connection for Virus Scan	Nblade.vscanNoScanner Conn	SVM	Critical
No AV Server Registered	Nblade.vscanNoRegdSca nner	SVM	Error
No Responsive AV Server Connection	Nblade.vscanConnInactiv e	SVM	Information
AV Server too Busy to Accept New Scan Request	Nblade.vscanConnBackPr essure	SVM	Error
Unauthorized User Attempt to AV Server	Nblade.vscanBadUserPriv Access	SVM	Error
FlexGroup Constituents Have Space Issues	flexgroup.constituents.hav e.space.issues	Volume	Error
FlexGroup Constituents Space Status All OK	flexgroup.constituents.spa ce.status.all.ok	Volume	Information
FlexGroup Constituents Have Inodes Issues	flexgroup.constituents.hav e.inodes.issues	Volume	Error
FlexGroup Constituents Inodes Status All OK	flexgroup.constituents.ino des.status.all.ok	Volume	Information
Volume Logical Space Nearly Full	monitor.vol.nearFull	Volume	Warning
Volume Logical Space Full	monitor.vol.full	Volume	Error
Volume Logical Space Normal	monitor.vol.one.ok	Volume	Information
WAFL Volume AutoSize Fail	wafl.vol.autoSize.fail	Volume	Error

Unified Manager Event name	EMS Event name	Affected resource	ONTAP severity
WAFL Volume AutoSize Done	wafl.vol.autoSize.done	Volume	Information

Subscribing to ONTAP EMS events

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

Before you begin

Do not subscribe to EMS events that are already added to Unified Manager automatically as this can cause confusion when receiving two events for the same issue.

About this task

You can subscribe to any number of EMS events. All the events to which you subscribe are validated, and only the validated events are applied to the clusters you are monitoring in Unified Manager. The *ONTAP 9 EMS Event Catalog* provides detailed information for all of the EMS messages for the specified version of ONTAP 9 software. Locate the appropriate version of the *EMS Event Catalog* from the ONTAP 9 Product Documentation page for a list of the applicable events.

ONTAP 9 Product Library

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



If you do not receive the ONTAP EMS events to which you have subscribed, there might be an issue with the DNS configuration of the cluster which is preventing the cluster from reaching the Unified Manager server. To resolve this issue, the cluster administrator must correct the DNS configuration of the cluster, and then restart Unified Manager. Doing so will flush the pending EMS events to the Unified Manager server.

Steps

- 1. In the left navigation pane, click **Configuration > Manage Events**.
- 2. In the Configuration/Manage Events page, click the Subscribe to EMS events button.
- 3. In the **Subscribe to EMS events** dialog box, enter the name of the ONTAP EMS event to which you want to subscribe.

To view the names of the EMS events to which you can subscribe, from the ONTAP cluster shell, you can use the event route show command (prior to ONTAP 9) or the event catalog show command (ONTAP 9 or later). See Knowledgebase Answer 1072320 for detailed instructions for identifying individual EMS events.

How to configure and receive alerts from ONTAP EMS Event Subscription in Active IQ Unified Manager

4. Click Add.

The EMS event is added to the Subscribed EMS events list, but the Applicable to Cluster column displays the status as "Unknown" for the EMS event that you added.

- 5. Click Save and Close to register the EMS event subscription with the cluster.
- 6. Click Subscribe to EMS events again.

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

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

After you finish

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

Viewing event details

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

Before you begin

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

About this task

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

Steps

- 1. In the left navigation pane, click **Events**.
- 2. In the **Events** inventory page, click the event name for which you want to view the details.

The event details are displayed in the Event details page.

Viewing unassigned events

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

Before you begin

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

Steps

- 1. In the left navigation pane, click **Events**.
 - By default, New and Acknowledged events are displayed on the Events inventory page.
- 2. From the Filters pane, select the Unassigned filter option in the Assigned To area.

Acknowledging and resolving events

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

Before you begin

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

About this task

You can acknowledge and resolve multiple events simultaneously.



You cannot acknowledge Information events.

Steps

- 1. In the left navigation pane, click **Events**.
- 2. From the events list, perform the following actions to acknowledge the events:

If you want to	Do this
Acknowledge and mark a single event as resolved	a. Click the event name.
	 b. From the Event details page, determine the cause of the event.
	c. Click Acknowledge .
	d. Take appropriate corrective action.
	e. Click Mark As Resolved.
Acknowledge and mark multiple events as resolved	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. In the Notes and Updates area, add a note about how you addressed the event, and then click Post.

Assigning events to specific users

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

Before you begin

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

Steps

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

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

Adding and reviewing notes about an event

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

Before you begin

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

Steps

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

Disabling or enabling events

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

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

Steps

- 1. In the left navigation pane, click Configuration > Manage Events.
- In the Configuration/Manage Events page page, disable or enable events by choosing one of the following options:

If you want to	Then do this
Disable events	a. Click Disable .
	 b. In the Disable Events dialog box, select the event severity.
	c. In the Matching Events column, select the events that you want to disable based on the event severity, and then click the right arrow to move those events to the Disable Events column.
	d. Click Save and Close.
	e. Verify that the events that you disabled are displayed in the list view of the Configuration/Manage Events page.

If you want to	Then do this
Enable events	a. Select the check box for the event, or events, that you want to enable.b. Click Enable.

What a Unified Manager maintenance window is

You define a Unified Manager maintenance window to suppress events and alerts for a specific timeframe when you have scheduled cluster maintenance and you do not want to receive a flood of unwanted notifications.

When the maintenance window starts, an "Object Maintenance Window Started" event is posted to the Events inventory page. This event is obsoleted automatically when the maintenance window ends.

During a maintenance window the events related to all objects on that cluster are still generated, but they do not appear in any of the UI pages, and no alerts or other types of notification are sent for these events. You can, however, view the events that were generated for all storage objects during a maintenance window by selecting one of the View options on the Events inventory page.

You can schedule a maintenance window to be initiated in the future, you can change the start and end times for a scheduled maintenance window, and you can cancel a scheduled maintenance window.

Scheduling a maintenance window to disable cluster event notifications

If you have a planned downtime for a cluster, for example, to upgrade the cluster or to move one of the nodes, you can suppress the events and alerts that would normally be generated during that timeframe by scheduling a Unified Manager maintenance window.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

During a maintenance window, the events related to all objects on that cluster are still generated, but they do not appear in the event page, and no alerts or other types of notification are sent for these events.

The time you enter for the maintenance window is based on the time at the Unified Manager server.

Steps

- 1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
- 2. In the Maintenance Mode column for the cluster, select the slider button and move it to the right.

The calendar window is displayed.

3. Select the start and end date and time for the maintenance window and click Apply.

The message "Scheduled" appears next to the slider button.

Results

When the start time is reached the cluster goes into maintenance mode and an "Object Maintenance Window Started" event is generated.

Changing or canceling a scheduled maintenance window

If you have configured a Unified Manager maintenance window to occur in the future, you can change the start and end times or cancel the maintenance window from occurring.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Canceling a currently running maintenance window is useful if you have completed cluster maintenance before the scheduled maintenance window end time and you want to start receiving events and alerts from the cluster again.

Steps

- 1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
- 2. In the Maintenance Mode column for the cluster:

If you want to	Perform this step
Change the timeframe for a scheduled maintenance window	a. Click the text "Scheduled" next to the slider button.
	b. Change the start and/or end date and time and click Apply .
Extend the length of an active maintenance window	a. Click the text "Active" next to the slider button.
	b. Change the end date and time and click Apply .
Cancel a scheduled maintenance window	Select the slider button and move it to the left.
Cancel an active maintenance window	Select the slider button and move it to the left.

Viewing events that occurred during a maintenance window

If necessary, you can view the events that were generated for all storage objects during a Unified Manager maintenance window. Most events will appear in the Obsolete state once the maintenance window has completed and all system resources are back up and running.

Before you begin

At least one maintenance window must have completed before any events are available.

About this task

Events that occurred during a maintenance window do not appear on the Events inventory page by default.

Steps

1. In the left navigation pane, click **Events**.

By default, all active (New and Acknowledged) events are displayed on the Events inventory page.

2. From the View pane, select the option All events generated during maintenance.

The list of events trigged during the last 7 days from all maintenance window sessions and from all clusters are displayed.

3. If there have been multiple maintenance windows for a single cluster, you can click the **Triggered Time** calendar icon and select the period of time for the maintenance window events that you are interested in viewing.

Managing host system resource events

Unified Manager includes a service that monitors resource issues on the host system on which Unified Manager is installed. Issues such as lack of available disk space or lack of memory on the host system may trigger management station events that are displayed as banner messages across the top of the UI.

About this task

Management station events indicate an issue with the host system on which Unified Manager is installed. Examples of management station issues include disk space running low on the host system; Unified Manager missing a regular data collection cycle; and noncompletion, or late completion, of statistics analysis because the next collection poll was initiated.

Unlike all other Unified Manager event messages, these particular management station warning and critical events are displayed in banner messages.

Steps

1. To view management station event information, perform these actions:

If you want to	Do this
View details of the event	Click the event banner to display the Event details page that includes suggested solutions for the issue.
View all management station events	a. In the left navigation pane, click Events.b. In the Filters pane on the Events inventory page, click the box for Management Station in the Source Type list.

Understanding more about events

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

Event state definitions

The state of an event helps you identify whether an appropriate corrective action is required. An event can be New, Acknowledged, Resolved, or Obsolete. Note that both New and Acknowledged events are considered to be active events.

The 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.



You cannot acknowledge or resolve an obsolete event.

Example of different states of an event

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

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

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

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

Description of event severity types

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

Critical

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

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

• Error

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

Warning

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

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

Information

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

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

Description of event impact levels

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

Incident

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

Risk

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

Event

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

Description of event impact areas

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

which you are responsible.

Availability

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

Capacity

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

Configuration

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

Performance

Performance events notify you of resource, configuration, or activity conditions on your cluster that might adversely affect the speed of data storage input or retrieval on your monitored storage objects.

Protection

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

How object status is computed

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

Sources of performance events

Performance events are issues related to workload performance on a cluster. They help you identify storage objects with slow response times, also known as high latency. Together with other health events that occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

Unified Manager receives performance events from the following sources:

User-defined performance threshold policy events

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

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

System-defined performance threshold policy events

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

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

Dynamic performance threshold events

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

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

Dynamic performance event chart details

For dynamic performance events, the System Diagnosis section of the Event details page lists the top workloads with the highest latency or usage of the cluster component that is in contention. The performance statistics are based on the time the performance event was detected up to the last time the event was analyzed. The charts also display historical performance statistics for the cluster component that is in contention.

For example, you can identify workloads with high utilization of a component to determine which workload to move to a less-utilized component. Moving the workload would reduce the amount of work on the current component, possibly bringing the component out of contention. At the of this section is the time and date range when an event was detected and last analyzed. For active events (new or acknowledged), the last analyzed time continues to update.

The latency and activity charts display the names of the top workloads when you hover your cursor over the chart. Clicking the Workload Type menu at the right of the chart enables you to sort the workloads based on their role in the event, including *sharks*, *bullies*, or *victims*, and displays details about their latency and their usage on the cluster component in contention. You can compare the actual value to the expected value to see when the workload was outside its expected range of latency or usage. See Workloads monitored by Unified Manager.



When you sort by peak deviation in latency, system-defined workloads are not displayed in the table, because latency applies only to user-defined workloads. Workloads with very low latency values are not displayed in the table.

For more information about the dynamic performance thresholds, see What events are. For information about how Unified Manager ranks the workloads and determines the sort order, see How Unified Manager determines the performance impact for an event.

The data in the graphs shows 24 hours of performance statistics prior to the last time the event was analyzed. The actual values and expected values for each workload are based on the time the workload was involved in the event. For example, a workload might become involved in an event after the event was detected, so its performance statistics might not match the values at the time of event detection. By default, the workloads are sorted by peak (highest) deviation in latency.



Because Unified Manager retains a maximum of 30 days of 5-minute historical performance and event data, if the event is more than 30 days old, no performance data is displayed.

· Workload Sort column

Latency chart

Displays the impact of the event to the latency of the workload during the last analysis.

Component Usage column

Displays details about the workload usage of the cluster component in contention. In the graphs, the actual usage is a blue line. A red bar highlights the event duration, from the detection time to the last analyzed time. For more information, see Workload performance measurements.



For the network component, because network performance statistics come from activity off the cluster, this column is not displayed.

Component Usage

Displays the history of utilization, in percent, for the network processing, data processing, and aggregate components or the history of activity, in percent, for the QoS policy group component. The chart is not displayed for the network or interconnect components. You can point to the statistics to view the usage statistics at a specific point in time.

Total Write MBps History

For the MetroCluster Resources component only, shows the total write throughput, in megabytes per second (MBps), for all volume workloads that are being mirrored to the partner cluster in a MetroCluster configuration.

Event History

Displays red-shaded lines to indicate the historic events for the component in contention. For obsolete events, the chart displays events that occurred before the selected event was detected and after it was resolved.

Types of system-defined performance threshold policies

Unified Manager provides some standard threshold policies that monitor cluster performance and generate events automatically. These policies are enabled by default, and they generate warning or information events when the monitored performance thresholds are breached.



System-defined performance threshold policies are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

If you are receiving unnecessary events from any system-defined performance threshold policies, you can disable individual policies from the Configuration/Manage Events page.

Node threshold policies

The system-defined node performance threshold policies are assigned, by default, to every node in the clusters being monitored by Unified Manager:

Node resources over-utilized

Identifies situations in which a single node is operating above the bounds of its operational efficiency, and therefore potentially affecting workload latencies. This is a warning event.

For nodes installed with ONTAP 8.3.x and earlier software, it does this by looking for nodes that are using more than 85% of their CPU and RAM resources (node utilization) for more than 30 minutes.

For nodes installed with ONTAP 9.0 and later software, it does this by looking for nodes that are using more than 100% of their performance capacity for more than 30 minutes.

Node HA pair over-utilized

Identifies situations in which nodes in an HA pair are operating above the bounds of the HA pair operational efficiency. This is an informational event.

For nodes installed with ONTAP 8.3.x and earlier software, it does this by looking at the CPU and RAM usage for the two nodes in the HA pair. If the combined node utilization of the two nodes exceeds 140% for more than one hour, then a controller failover will impact workload latencies.

For nodes installed with ONTAP 9.0 and later software, it does this by looking at the performance capacity used value for the two nodes in the HA pair. If the combined performance capacity used of the two nodes exceeds 200% for more than one hour, then a controller failover will impact workload latencies.

Node disk fragmentation

Identifies situations in which a disk or disks in an aggregate are fragmented, slowing key system services and potentially affecting workload latencies on a node.

It does this by looking at certain read and write operation ratios across all aggregates on a node. This policy might also be triggered during SyncMirror resynchronization or when errors are found during disk scrub operations. This is a warning event.



The "Node disk fragmentation" policy analyzes HDD-only aggregates; Flash Pool, SSD, and FabricPool aggregates are not analyzed.

Aggregate threshold policies

The system-defined aggregate performance threshold policy is assigned by default to every aggregate in the clusters being monitored by Unified Manager.

Aggregate disks over-utilized

Identifies situations in which an aggregate is operating above the limits of its operational efficiency, thereby potentially affecting workload latencies. It identifies these situations by looking for aggregates where the disks in the aggregate are more than 95% utilized for more than 30 minutes. This multicondition policy then performs the following analysis to help determine the cause of the issue:

• Is a disk in the aggregate currently undergoing background maintenance activity?

Some of the background maintenance activities a disk could be undergoing are disk reconstruction, disk scrub, SyncMirror resynchronization, and reparity.

- Is there a communications bottleneck in the disk shelf Fibre Channel interconnect?
- Is there too little free space in the aggregate?
 A warning event is issued for this policy only if one (or more) of the three subordinate policies are also considered breached. A performance event is not triggered if only the disks in the aggregate are more than 95% utilized.



The "Aggregate disks over-utilized" policy analyzes HDD-only aggregates and Flash Pool (hybrid) aggregates; SSD and FabricPool aggregates are not analyzed.

QoS threshold policies

The system-defined QoS performance threshold policies are assigned to any workload that has a configured ONTAP QoS maximum throughput policy (IOPS, IOPS/TB, or MBps). Unified Manager triggers an event when the workload throughput value is 15% less than the configured QoS value.

QoS Max IOPS or MBps threshold

Identifies volumes and LUNs that have exceeded their QoS maximum IOPS or MBps throughput limit, and that are affecting workload latency. This is a warning event.

When a single workload is assigned to a policy group, it does this by looking for workloads that have exceeded the maximum throughput threshold defined in the assigned QoS policy group during each collection period for the previous hour.

When multiple workloads share a single QoS policy, it does this by adding the IOPS or MBps of all workloads in the policy and checking that total against the threshold.

• QoS Peak IOPS/TB or IOPS/TB with Block Size threshold

Identifies volumes that have exceeded their adaptive QoS peak IOPS/TB throughput limit (or IOPS/TB with Block Size limit), and that are affecting workload latency. This is a warning event.

It does this by converting the peak IOPS/TB threshold defined in the adaptive QoS policy into a QoS maximum IOPS value based on the size of each volume, and then it looks for volumes that have exceeded the QoS max IOPS during each performance collection period for the previous hour.



This policy is applied to volumes only when the cluster is installed with ONTAP 9.3 and later software.

When the "block size" element has been defined in the adaptive QoS policy, the threshold is converted into a QoS maximum MBps value based on the size of each volume. Then it looks for volumes that have exceeded the QoS max MBps during each performance collection period for the previous hour.



This policy is applied to volumes only when the cluster is installed with ONTAP 9.5 and later software.

List of events and severity types

You can use the list of events to become more familiar with event categories, event

names, and the severity type of each event that you might see in Unified Manager. Events are listed in alphabetical order by object category.

Aggregate events

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

Impact area: availability

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name(Trap name)	Impact level	Source type	Severity
Aggregate Offline(ocumEvtAggregat eStateOffline)	Incident	Aggregate	Critical
Aggregate Failed(ocumEvtAggregate StateFailed)	Incident	Aggregate	Critical
Aggregate Restricted(ocumEvtAggre gateStateRestricted)	Risk	Aggregate	Warning
Aggregate Reconstructing(ocumEvtA ggregateRaidStateRecon structing)	Risk	Aggregate	Warning
Aggregate Degraded(ocumEvtAggre gateRaidStateDegraded)	Risk	Aggregate	Warning
Cloud Tier Partially Reachable(ocumEventClo udTierPartiallyReachable)	Risk	Aggregate	Warning
Cloud Tier Unreachable(ocumEvent CloudTierUnreachable)	Risk	Aggregate	Error
MetroCluster Aggregate Left Behind(ocumEvtMetroClu sterAggregateLeftBehind)	Risk	Aggregate	Error

Event name(Trap name)	Impact level	Source type	Severity
MetroCluster Aggregate Mirroring Degraded(ocumEvtMetro ClusterAggregateMirrorDe graded)	Risk	Aggregate	Error
Object-store Access Denied for Aggregate Relocation *	Risk	Aggregate	Error
Object-store Access Denied for Aggregate Relocation During Storage Failover *	Risk	Aggregate	Error

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
Aggregate Space Nearly Full(ocumEvtAggregateN earlyFull)	Risk	Aggregate	Warning
Aggregate Space Full(ocumEvtAggregateFu II)	Risk	Aggregate	Error
Aggregate Days Until Full(ocumEvtAggregateD aysUntilFullSoon)	Risk	Aggregate	Error
Aggregate Overcommitted(ocumEvtA ggregateOvercommitted)	Risk	Aggregate	Error
Aggregate Nearly Overcommitted(ocumEvtA ggregateAlmostOvercom mitted)	Risk	Aggregate	Warning
Aggregate Snapshot Reserve Full(ocumEvtAggregateSn apReserveFull)	Risk	Aggregate	Warning

Event name(Trap name)	Impact level	Source type	Severity
Aggregate Growth Rate Abnormal(ocumEvtAggre gateGrowthRateAbnormal)	Risk	Aggregate	Warning

Impact area: configuration

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

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
Aggregate IOPS Critical Threshold Breached(ocumAggregate IopsIncident)	Incident	Aggregate	Critical
Aggregate IOPS Warning Threshold Breached(ocumAggregate IopsWarning)	Risk	Aggregate	Warning
Aggregate MBps Critical Threshold Breached(ocumAggregate MbpsIncident)	Incident	Aggregate	Critical
Aggregate MBps Warning Threshold Breached(ocumAggregateMbpsWar ning)	Risk	Aggregate	Warning
Aggregate Latency Critical Threshold Breached(ocumAggregate LatencyIncident)	Incident	Aggregate	Critical

Event name(Trap name)	Impact level	Source type	Severity
Aggregate Latency Warning Threshold Breached(ocumAggregateLatencyW arning)	Risk	Aggregate	Warning
Aggregate Perf. Capacity Used Critical Threshold Breached(ocumAggregate PerfCapacityUsedIncident)	Incident	Aggregate	Critical
Aggregate Perf. Capacity Used Warning Threshold Breached(ocumAggregate PerfCapacityUsedWarnin g)	Risk	Aggregate	Warning
Aggregate Utilization Critical Threshold Breached (ocumAggregateUtilizatio nIncident)	Incident	Aggregate	Critical
Aggregate Utilization Warning Threshold Breached (ocumAggregateUtilizatio nWarning)	Risk	Aggregate	Warning
Aggregate Disks Over- utilized Threshold Breached (ocumAggregateDisksOve rUtilizedWarning)	Risk	Aggregate	Warning
Aggregate Dynamic Threshold Breached (ocumAggregateDynamic EventWarning)	Risk	Aggregate	Warning

Cluster events

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

Impact area: availability

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name(Trap name)	Impact level	Source type	Severity
Cluster Lacks Spare Disks(ocumEvtDisksNoSp ares)	Risk	Cluster	Warning
Cluster Not Reachable(ocumEvtClust erUnreachable)	Risk	Cluster	Error
Cluster Monitoring Failed(ocumEvtClusterMo nitoringFailed)	Risk	Cluster	Warning
Cluster FabricPool License Capacity Limits Breached (ocumEvtExternalCapacit yTierSpaceFull)	Risk	Cluster	Warning
NVMe-oF Grace Period Started *(nvmfGracePeriodStart)	Risk	Cluster	Warning
NVMe-oF Grace Period Active *(nvmfGracePeriodActive)	Risk	Cluster	Warning
NVMe-oF Grace Period Expired *(nvmfGracePeriodExpire d)	Risk	Cluster	Warning
Object Maintenance Window Started(objectMaintenanc eWindowStarted)	Event	Cluster	Critical
Object Maintenance Window Ended(objectMaintenance WindowEnded)	Event	Cluster	Information

Event name(Trap name)	Impact level	Source type	Severity
MetroCluster Spare Disks Left Behind(ocumEvtSpareDis kLeftBehind)	Risk	Cluster	Error
MetroCluster Automatic Unplanned Switchover Disabled(ocumEvtMccAut omaticUnplannedSwitchO verDisabled)	Risk	Cluster	Warning

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
Cluster Cloud Tier Planning (clusterCloudTierPlanning Warning)	Risk	Cluster	Warning
FabricPool Space Nearly Full *	Risk	Cluster	Error

Impact area: configuration

Event name(Trap name)	Impact level	Source type	Severity
Node Added(Not applicable)	Event	Cluster	Information
Node Removed(Not applicable)	Event	Cluster	Information
Cluster Removed(Not applicable)	Event	Cluster	Information
Cluster Add Failed(Not applicable)	Event	Cluster	Error
Cluster Name Changed(Not applicable)	Event	Cluster	Information
Emergency EMS received (Not applicable)	Event	Cluster	Critical

Event name(Trap name)	Impact level	Source type	Severity
Critical EMS received (Not applicable)	Event	Cluster	Critical
Alert EMS received (Not applicable)	Event	Cluster	Error
Error EMS received (Not applicable)	Event	Cluster	Warning
Warning EMS received (Not applicable)	Event	Cluster	Warning
Debug EMS received (Not applicable)	Event	Cluster	Warning
Notice EMS received (Not applicable)	Event	Cluster	Warning
Informational EMS received (Not applicable)	Event	Cluster	Warning

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

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

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
Cluster IOPS Critical Threshold Breached(ocumClusterlop sIncident)	Incident	Cluster	Critical
Cluster IOPS Warning Threshold Breached(ocumClusterlop sWarning)	Risk	Cluster	Warning
Cluster MBps Critical Threshold Breached(ocumClusterMb psIncident)	Incident	Cluster	Critical
Cluster MBps Warning Threshold Breached(ocumClusterMb psWarning)	Risk	Cluster	Warning
Cluster Dynamic Threshold Breached(ocumClusterDy namicEventWarning)	Risk	Cluster	Warning

Disks events

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

Event name(Trap name)	Impact level	Source type	Severity
Flash Disks - Spare Blocks Almost Consumed(ocumEvtClust erFlashDiskFewerSpareBl ockError)	Risk	Cluster	Error
Flash Disks - No Spare Blocks(ocumEvtClusterFl ashDiskNoSpareBlockCrit ical)	Incident	Cluster	Critical

Event name(Trap name)	Impact level	Source type	Severity
Some Unassigned Disks(ocumEvtClusterUna ssignedDisksSome)	Risk	Cluster	Warning
Some Failed Disks(ocumEvtDisksSom eFailed)	Incident	Cluster	Critical

Enclosures events

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

Event name(Trap name)	Impact level	Source type	Severity
Disk Shelf Fans Failed(ocumEvtShelfFanF ailed)	Incident	Storage shelf	Critical
Disk Shelf Power Supplies Failed(ocumEvtShelfPow erSupplyFailed)	Incident	Storage shelf	Critical
Disk Shelf Multipath Not Configured(ocumDiskShel fConnectivityNotInMultiPa th) This event does not apply to: Clusters that are in a MetroCluster configuration The following platforms: FAS2554, FAS2552, FAS2520, and FAS2240	Risk	Node	Warning
Disk Shelf Path Failure(ocumDiskShelfCo nnectivityPathFailure)	Risk	Storage Shelf	Warning

Impact area: configuration

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

Fans events

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

Impact area: availability

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

Flash card events

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

Impact area: availability

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

Inodes events

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

Impact area: capacity

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

Logical interface (LIF) events

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

Impact area: availability

Event name(Trap name)	Impact level	Source type	Severity
LIF Status Down(ocumEvtLifStatusD own)	Risk	Interface	Error
LIF Failover Not Possible(ocumEvtLifFailo verNotPossible)	Risk	Interface	Warning
LIF Not At Home Port(ocumEvtLifNotAtHo mePort)	Risk	Interface	Warning

Impact area: configuration

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

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
Network LIF MBps Critical Threshold Breached(ocumNetworkLi fMbpsIncident)		Interface	Critical

Event name(Trap name)	Impact level	Source type	Severity
Network LIF MBps Warning Threshold Breached(ocumNetworkLifMbpsWarning)	Risk	Interface	Warning
FCP LIF MBps Critical Threshold Breached(ocumFcpLifMb psIncident)	Incident	Interface	Critical
FCP LIF MBps Warning Threshold Breached(ocumFcpLifMb psWarning)	Risk	Interface	Warning
NVMf FCP LIF MBps Critical Threshold Breached(ocumNvmfFcLif MbpsIncident)	Incident	Interface	Critical
NVMf FCP LIF MBps Warning Threshold Breached(ocumNvmfFcLif MbpsWarning)	Risk	Interface	Warning

LUN events

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

Impact area: availability

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

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

Event name(Trap name)	Impact level	Source type	Severity
No Active Paths To Access LUN(ocumEvtLunNotRea chable)	Incident	LUN	Critical
No Optimized Paths To Access LUN(ocumEvtLunOptimiz edPathInactive)	Risk	LUN	Warning
No Paths To Access LUN From HA Partner(ocumEvtLunHaPa thInactive)	Risk	LUN	Warning

Impact area: capacity

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

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
LUN IOPS Critical Threshold Breached(ocumLunlopsIn cident)	Incident	LUN	Critical
LUN IOPS Warning Threshold Breached(ocumLunlopsW arning)	Risk	LUN	Warning
LUN MBps Critical Threshold Breached(ocumLunMbpsl ncident)	Incident	LUN	Critical
LUN MBps Warning Threshold Breached(ocumLunMbps Warning)	Risk	LUN	Warning

Event name(Trap name)	Impact level	Source type	Severity
LUN Latency ms/op Critical Threshold Breached(ocumLunLaten cylncident)	Incident	LUN	Critical
LUN Latency ms/op Warning Threshold Breached(ocumLunLaten cyWarning)	Risk	LUN	Warning
LUN Latency and IOPS Critical Threshold Breached(ocumLunLaten cylopsIncident)	Incident	LUN	Critical
LUN Latency and IOPS Warning Threshold Breached(ocumLunLaten cylopsWarning)	Risk	LUN	Warning
LUN Latency and MBps Critical Threshold Breached(ocumLunLaten cyMbpsIncident)	Incident	LUN	Critical
LUN Latency and MBps Warning Threshold Breached(ocumLunLaten cyMbpsWarning)	Risk	LUN	Warning
LUN Latency and Aggregate Perf. Capacity Used Critical Threshold Breached(ocumLunLaten cyAggregatePerfCapacity UsedIncident)	Incident	LUN	Critical
LUN Latency and Aggregate Perf. Capacity Used Warning Threshold Breached(ocumLunLaten cyAggregatePerfCapacity UsedWarning)	Risk	LUN	Warning

Event name(Trap name)	Impact level	Source type	Severity
LUN Latency and Aggregate Utilization Critical Threshold Breached(ocumLunLaten cyAggregateUtilizationInci dent)	Incident	LUN	Critical
LUN Latency and Aggregate Utilization Warning Threshold Breached(ocumLunLaten cyAggregateUtilizationWa rning)	Risk	LUN	Warning
LUN Latency and Node Perf. Capacity Used Critical Threshold Breached(ocumLunLaten cyNodePerfCapacityUsed Incident)	Incident	LUN	Critical
LUN Latency and Node Perf. Capacity Used Warning Threshold Breached(ocumLunLaten cyNodePerfCapacityUsed Warning)	Risk	LUN	Warning
LUN Latency and Node Perf. Capacity Used - Takeover Critical Threshold Breached(ocumLunLaten cyAggregatePerfCapacity UsedTakeoverIncident)	Incident	LUN	Critical
LUN Latency and Node Perf. Capacity Used - Takeover Warning Threshold Breached(ocumLunLaten cyAggregatePerfCapacity UsedTakeoverWarning)	Risk	LUN	Warning
LUN Latency and Node Utilization Critical Threshold Breached(ocumLunLaten cyNodeUtilizationIncident)	Incident	LUN	Critical

Event name(Trap name)	Impact level	Source type	Severity
LUN Latency and Node Utilization Warning Threshold Breached(ocumLunLaten cyNodeUtilizationWarning)	Risk	LUN	Warning
QoS LUN Max IOPS Warning Threshold Breached(ocumQosLunM axlopsWarning)	Risk	LUN	Warning
QoS LUN Max MBps Warning Threshold Breached(ocumQosLunM axMbpsWarning)	Risk	LUN	Warning

Management station events

Management station events provide you with information about the status of server on which Unified Manager is installed so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: configuration

Event name(Trap name)	Impact level	Source type	Severity
Unified Manager Server Disk Space Nearly Full(ocumEvtUnifiedMana gerDiskSpaceNearlyFull)	Risk	Management station	Warning
Unified Manager Server Disk Space Full(ocumEvtUnifiedMana gerDiskSpaceFull)	Incident	Management station	Critical
Unified Manager Server Low On Memory(ocumEvtUnified ManagerMemoryLow)	Risk	Management station	Warning
Unified Manager Server Almost Out Of Memory(ocumEvtUnified ManagerMemoryAlmostO ut)	Incident	Management station	Critical

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
Performance Data Analysis Is Impacted(ocumEvtUnified ManagerDataMissingAnal yze)	Risk	Management station	Warning
Performance Data Collection Is Impacted(ocumEvtUnified ManagerDataMissingColl ection)	Incident	Management station	Critical



These last two performance events were available for Unified Manager 7.2 only. If either of these events exist in the New state, and then you upgrade to a newer version of Unified Manager software, the events will not be purged automatically. You will need to move the events to the Resolved state manually.

MetroCluster Bridge events

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

Impact area: availability

Event name(Trap name)	Impact level	Source type	Severity
Bridge Unreachable(ocumEvtBrid geUnreachable)	Incident	MetroCluster Bridge	Critical
Bridge Temperature Abnormal(ocumEvtBridge TemperatureAbnormal)	Incident	MetroCluster Bridge	Critical

MetroCluster Connectivity events

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

Event name(Trap name)	Impact level	Source type	Severity
All Inter-Switch Links Down(ocumEvtMetroClust erAllISLBetweenSwitches Down)	Incident	MetroCluster inter-switch connection	Critical
All Links Between MetroCluster Partners Down(ocumEvtMetroClust erAllLinksBetweenPartner sDown)	Incident	MetroCluster relationship	Critical
FC-SAS Bridge To Storage Stack Link Down(ocumEvtBridgeSas PortDown)	Incident	MetroCluster bridge stack connection	Critical
MetroCluster Configuration Switched Over((ocumEvtMetroClust erDRStatusImpacted)	Risk	MetroCluster relationship	Warning
MetroCluster Configuration Partially Switched Over(ocumEvtMetroClust erDRStatusPartiallyImpac ted)	Risk	MetroCluster relationship	Error
MetroCluster Disaster Recovery Capability Impacted(ocumEvtMetroC lusterDRStatusImpacted)	Risk	MetroCluster relationship	Critical
MetroCluster Partners Not Reachable Over Peering Network(ocumEvtMetroCl usterPartnersNotReachab leOverPeeringNetwork)	Incident	MetroCluster relationship	Critical
Node To FC Switch All FC-VI Interconnect Links Down(ocumEvtMccNodeS witchFcviLinksDown)	Incident	MetroCluster node switch connection	Critical

Event name(Trap name)	Impact level	Source type	Severity
Node To FC Switch One Or More FC-Initiator Links Down(ocumEvtMccNodeS witchFcLinksOneOrMore Down)	Risk	MetroCluster node switch connection	Warning
Node To FC Switch All FC-Initiator Links Down(ocumEvtMccNodeS witchFcLinksDown)	Incident	MetroCluster node switch connection	Critical
Switch To FC-SAS Bridge FC Link Down (ocumEvtMccSwitchBridg eFcLinksDown)	Incident	MetroCluster switch bridge connection	Critical
Inter Node All FC VI InterConnect Links Down (ocumEvtMccInterNodeLi nksDown)	Incident	Inter-node connection	Critical
Inter Node One Or More FC VI InterConnect Links Down (ocumEvtMccInterNodeLi nksOneOrMoreDown)	Risk	Inter-node connection	Warning
Node To Bridge Link Down (ocumEvtMccNodeBridge LinksDown)	Incident	Node bridge connection	Critical
Node to Storage Stack All SAS Links Down (ocumEvtMccNodeStackLi nksDown)	Incident	Node stack connection	Critical
Node to Storage Stack One Or More SAS Links Down (ocumEvtMccNodeStackLi nksOneOrMoreDown)	Risk	Node stack connection	Warning

MetroCluster switch events

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

and severity.

Impact area: availability

Event na	me(Trap name)	Impact level	Source type	Severity
Switch Temperature Abnormal(ocumEvtSwitch TemperatureAbnormal)		Incident	MetroCluster Switch	Critical
Switch Unreachable(ocumEvtSwitchUnreachable)		Incident	MetroCluster Switch	Critical
Switch Fans Failed(ocumEvtSwitchFan sOneOrMoreFailed)		Incident	MetroCluster Switch	Critical
Switch Power Supplies Failed(ocumEvtSwitchPo werSuppliesOneOrMoreF ailed)		Incident	MetroCluster Switch	Critical
Switch Temperature Sensors Failed(ocumEvtSwitchTe mperatureSensorFailed)		Incident	MetroCluster Switch	Critical
i	This event is applicable only for Cisco switches.			

NVMe Namespace events

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

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name(Trap name)	Impact level	Source type	Severity
NVMeNS Offline *(nvmeNamespaceStatus Offline)	Event	Namespace	Information

Event name(Trap name)	Impact level	Source type	Severity
NVMeNS Online *(nvmeNamespaceStatus Online)	Event	Namespace	Information
NVMeNS Out of Space *(nvmeNamespaceSpace OutOfSpace)	Risk	Namespace	Warning
NVMeNS Destroy *(nvmeNamespaceDestro y)	Event	Namespace	Information

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
NVMe Namespace IOPS Critical Threshold Breached(ocumNvmeNa mespacelopsIncident)	Incident	Namespace	Critical
NVMe Namespace IOPS Warning Threshold Breached(ocumNvmeNa mespacelopsWarning)	Risk	Namespace	Warning
NVMe Namespace MBps Critical Threshold Breached(ocumNvmeNa mespaceMbpsIncident)	Incident	Namespace	Critical
NVMe Namespace MBps Warning Threshold Breached(ocumNvmeNa mespaceMbpsWarning)	Risk	Namespace	Warning
NVMe Namespace Latency ms/op Critical Threshold Breached(ocumNvmeNa mespaceLatencyIncident)	Incident	Namespace	Critical
NVMe Namespace Latency ms/op Warning Threshold Breached(ocumNvmeNa mespaceLatencyWarning)	Risk	Namespace	Warning

Event name(Trap name)	Impact level	Source type	Severity
NVMe Namespace Latency and IOPS Critical Threshold Breached(ocumNvmeNa mespaceLatencylopsIncid ent)	Incident	Namespace	Critical
NVMe Namespace Latency and IOPS Warning Threshold Breached(ocumNvmeNa mespaceLatencylopsWar ning)	Risk	Namespace	Warning
NVMe Namespace Latency and MBps Critical Threshold Breached(ocumNvmeNa mespaceLatencyMbpsInci dent)	Incident	Namespace	Critical
NVMe Namespace Latency and MBps Warning Threshold Breached(ocumNvmeNa mespaceLatencyMbpsWa rning)	Risk	Namespace	Warning

Node events

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

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

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

Event name(Trap name)	Impact level	Source type	Severity
Cloud AWS MetaDataConnFail *(ocumCloudAwsMetadat aConnFail)	Risk	Node	Error
Cloud AWS IAMCredsExpired *(ocumCloudAwslamCred sExpired)	Risk	Node	Error
Cloud AWS IAMCredsInvalid *(ocumCloudAwslamCred sInvalid)	Risk	Node	Error
Cloud AWS IAMCredsNotFound *(ocumCloudAwslamCred sNotFound)	Risk	Node	Error
Cloud AWS IAMCredsNotInitialized *(ocumCloudAwslamCred sNotInitialized)	Event	Node	Information
Cloud AWS IAMRoleInvalid *(ocumCloudAwslamRoleI nvalid)	Risk	Node	Error
Cloud AWS IAMRoleNotFound *(ocumCloudAwslamRole NotFound)	Risk	Node	Error
Objstore Host Unresolvable *(ocumObjstoreHostUnres olvable)	Risk	Node	Error
Objstore InterClusterLifDown *(ocumObjstoreInterClust erLifDown)	Risk	Node	Error
Request Mismatch Object-store Signature *	Risk	Node	Error

Event name(Trap name)	Impact level	Source type	Severity
One of NFSv4 Pools Exhausted *	Incident	Node	Critical

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
QoS Monitor Memory Maxed *(ocumQosMonitorMemor yMaxed)	Risk	Node	Error
QoS Monitor Memory Abated *(ocumQosMonitorMemor yAbated)	Event	Node	Information

Impact area: configuration

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

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
Node IOPS Critical Threshold Breached(ocumNodelopsI ncident)	Incident	Node	Critical
Node IOPS Warning Threshold Breached(ocumNodelops Warning)	Risk	Node	Warning
Node MBps Critical Threshold Breached(ocumNodeMbp sIncident)	Incident	Node	Critical
Node MBps Warning Threshold Breached(ocumNodeMbp sWarning)	Risk	Node	Warning

Event name(Trap name)	Impact level	Source type	Severity
Node Latency ms/op Critical Threshold Breached(ocumNodeLate ncyIncident)	Incident	Node	Critical
Node Latency ms/op Warning Threshold Breached(ocumNodeLate ncyWarning)	Risk	Node	Warning
Node Perf. Capacity Used Critical Threshold Breached(ocumNodePerf CapacityUsedIncident)	Incident	Node	Critical
Node Perf. Capacity Used Warning Threshold Breached(ocumNodePerf CapacityUsedWarning)	Risk	Node	Warning
Node Perf.Capacity Used - Takeover Critical Threshold Breached(ocumNodePerf CapacityUsedTakeoverIn cident)	Incident	Node	Critical
Node Perf.Capacity Used - Takeover Warning Threshold Breached(ocumNodePerf CapacityUsedTakeoverW arning)	Risk	Node	Warning
Node Utilization Critical Threshold Breached (ocumNodeUtilizationIncid ent)	Incident	Node	Critical
Node Utilization Warning Threshold Breached (ocumNodeUtilizationWarning)	Risk	Node	Warning

Event name(Trap name)	Impact level	Source type	Severity
Node HA Pair Over- utilized Threshold Breached (ocumNodeHaPairOverUti lizedInformation)	Event	Node	Information
Node Disk Fragmentation Threshold Breached (ocumNodeDiskFragment ationWarning)	Risk	Node	Warning
Node Over-utilized Threshold Breached (ocumNodeOverUtilizedW arning)	Risk	Node	Warning
Node Dynamic Threshold Breached (ocumNodeDynamicEvent Warning)	Risk	Node	Warning

NVRAM battery events

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

Impact area: availability

Event name(Trap name)	Impact level	Source type	Severity
NVRAM Battery Low(ocumEvtNvramBatter yLow)	Risk	Node	Warning
NVRAM Battery Discharged(ocumEvtNvra mBatteryDischarged)	Risk	Node	Error
NVRAM Battery Overly Charged(ocumEvtNvramB atteryOverCharged)	Incident	Node	Critical

Port events

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

Impact area: availability

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

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
Network Port MBps Critical Threshold Breached(ocumNetworkP ortMbpsIncident)	Incident	Port	Critical
Network Port MBps Warning Threshold Breached(ocumNetworkP ortMbpsWarning)	Risk	Port	Warning
FCP Port MBps Critical Threshold Breached(ocumFcpPortM bpsIncident)	Incident	Port	Critical
FCP Port MBps Warning Threshold Breached(ocumFcpPortM bpsWarning)	Risk	Port	Warning
Network Port Utilization Critical Threshold Breached(ocumNetworkP ortUtilizationIncident)	Incident	Port	Critical
Network Port Utilization Warning Threshold Breached(ocumNetworkP ortUtilizationWarning)	Risk	Port	Warning
FCP Port Utilization Critical Threshold Breached(ocumFcpPortUt ilizationIncident)	Incident	Port	Critical

Event name(Trap name)	Impact level	Source type	Severity
FCP Port Utilization Warning Threshold Breached(ocumFcpPortUt ilizationWarning)	Risk	Port	Warning

Power supplies events

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

Impact area: availability

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

Protection events

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

Impact area: protection

Event name(Trap name)	Impact level	Source type	Severity
Protection Job Failed(ocumEvtProtection JobTaskFailed)	Incident	Volume or storage service	Critical
Protection Job Aborted(ocumEvtProtectio nJobAborted)	Risk	Volume or storage service	Warning

Qtree events

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

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
Qtree Space Nearly Full(ocumEvtQtreeSpace NearlyFull)	Risk	Qtree	Warning
Qtree Space Full(ocumEvtQtreeSpace Full)	Risk	Qtree	Error
Qtree Space Normal(ocumEvtQtreeSp aceThresholdOk)	Event	Qtree	Information
Qtree Files Hard Limit Reached(ocumEvtQtreeFi lesHardLimitReached)	Incident	Qtree	Critical
Qtree Files Soft Limit Breached(ocumEvtQtreeF ilesSoftLimitBreached)	Risk	Qtree	Warning
Qtree Space Hard Limit Reached(ocumEvtQtreeS paceHardLimitReached)	Incident	Qtree	Critical
Qtree Space Soft Limit Breached(ocumEvtQtreeS paceSoftLimitBreached)	Risk	Qtree	Warning

Service processor events

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

Event name(Trap name)	Impact level	Source type	Severity
Service Processor Not Configured(ocumEvtServi ceProcessorNotConfigure d)	Risk	Node	Warning
Service Processor Offline(ocumEvtServicePr ocessorOffline)	Risk	Node	Error

SnapMirror relationship events

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

Impact area: protection

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name(Trap name)	Impact level	Source type	Severity
Mirror Replication Unhealthy(ocumEvtSnap mirrorRelationshipUnhealt hy)	Risk	SnapMirror relationship	Warning
Mirror Replication Broken- off(ocumEvtSnapmirrorRe lationshipStateBrokenoff)	Risk	SnapMirror relationship	Error
Mirror Replication Initialize Failed(ocumEvtSnapmirro rRelationshipInitializeFaile d)	Risk	SnapMirror relationship	Error
Mirror Replication Update Failed(ocumEvtSnapmirro rRelationshipUpdateFaile d)	Risk	SnapMirror relationship	Error
Mirror Replication Lag Error(ocumEvtSnapMirror RelationshipLagError)	Risk	SnapMirror relationship	Error
Mirror Replication Lag Warning(ocumEvtSnapMir rorRelationshipLagWarnin g)	Risk	SnapMirror relationship	Warning
Mirror Replication Resync Failed(ocumEvtSnapmirro rRelationshipResyncFaile d)	Risk	SnapMirror relationship	Error
Mirror Replication DeletedocumEvtSnapmirr orRelationshipDeleted	Risk	SnapMirror relationship	Warning

Event name(Trap name)	Impact level	Source type	Severity
Synchronous Replication Out Of Sync *	Risk	SnapMirror relationship	Warning
Synchronous Replication Restored *	Event	SnapMirror relationship	Information
Synchronous Replication Auto Resync Failed *	Risk	SnapMirror relationship	Error

Snapshot events

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

Impact area: availability

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

SnapVault relationship events

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

Impact area: protection

Event name(Trap name)	Impact level	Source type	Severity
Asynchronous Vault Unhealthy(ocumEvtSnap VaultRelationshipUnhealt hy)	Risk	SnapMirror relationship	Warning

Event name(Trap name)	Impact level	Source type	Severity
Asynchronous Vault Broken- off(ocumEvtSnapVaultRel ationshipStateBrokenoff)	Risk	SnapMirror relationship	Error
Asynchronous Vault Initialize Failed(ocumEvtSnapVault RelationshipInitializeFaile d)	Risk	SnapMirror relationship	Error
Asynchronous Vault Update Failed(ocumEvtSnapVault RelationshipUpdateFailed)	Risk	SnapMirror relationship	Error
Asynchronous Vault Lag Error(ocumEvtSnapVault RelationshipLagError)	Risk	SnapMirror relationship	Error
Asynchronous Vault Lag Warning(ocumEvtSnapVa ultRelationshipLagWarnin g)	Risk	SnapMirror relationship	Warning
Asynchronous Vault Resync Failed(ocumEvtSnapvault RelationshipResyncFailed)	Risk	SnapMirror relationship	Error

Storage failover settings events

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

Event name(Trap name)	Impact level	Source type	Severity
Storage Failover Interconnect One Or More Links Down(ocumEvtSfoInterco nnectOneOrMoreLinksDo wn)	Risk	Node	Warning
Storage Failover Disabled(ocumEvtSfoSetti ngsDisabled)	Risk	Node	Error
Storage Failover Not Configured(ocumEvtSfoS ettingsNotConfigured)	Risk	Node	Error
Storage Failover State - Takeover(ocumEvtSfoStat eTakeover)	Risk	Node	Warning
Storage Failover State - Partial Giveback(ocumEvtSfoStat ePartialGiveback)	Risk	Node	Error
Storage Failover Node Status Down(ocumEvtSfoNodeSt atusDown)	Risk	Node	Error
Storage Failover Takeover Not Possible(ocumEvtSfoTak eoverNotPossible)	Risk	Node	Error

Storage services events

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

Impact area: configuration

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

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

Impact area: protection

Event name(Trap name)	Impact level	Source type	Severity
Unexpected Deletion of Managed SnapMirror RelationshipocumEvtStor ageServiceUnsupportedR elationshipDeletion	Risk	Storage service	Warning
Unexpected Deletion of Storage Service Member Volume(ocumEvtStorage ServiceUnexpectedVolum eDeletion)	Incident	Storage service	Critical

Storage shelf events

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

Event name(Trap name)	Impact level	Source type	Severity
Abnormal Voltage Range(ocumEvtShelfVolt ageAbnormal)	Risk	Storage shelf	Warning
Abnormal Current Range(ocumEvtShelfCurr entAbnormal)	Risk	Storage shelf	Warning
Abnormal Temperature(ocumEvtSh elfTemperatureAbnormal)	Risk	Storage shelf	Warning

SVM events

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

Impact area: availability

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name(Trap name)	Impact level	Source type	Severity
SVM CIFS Service Down(ocumEvtVserverCif sServiceStatusDown)	Incident	SVM	Critical
SVM CIFS Service Not Configured(Not applicable)	Event	SVM	Information
Attempts to Connect Nonexistent CIFS Share *	Incident	SVM	Critical
CIFS NetBIOS Name Conflict *	Risk	SVM	Error
CIFS Shadow Copy Operation Failed *	Risk	SVM	Error
Many CIFS Connections *	Risk	SVM	Error
Max CIFS Connection Exceeded *	Risk	SVM	Error
Max Number of CIFS Connection Per User Exceeded *	Risk	SVM	Error
SVM FC/FCoE Service Down(ocumEvtVserverFc ServiceStatusDown)	Incident	SVM	Critical
SVM iSCSI Service Down(ocumEvtVserverIsc siServiceStatusDown)	Incident	SVM	Critical
SVM NFS Service Down(ocumEvtVserverNf sServiceStatusDown)	Incident	SVM	Critical

Event name(Trap name)	Impact level	Source type	Severity
SVM FC/FCoE Service Not Configured(Not applicable)	Event	SVM	Information
SVM iSCSI Service Not Configured(Not applicable)	Event	SVM	Information
SVM NFS Service Not Configured(Not applicable)	Event	SVM	Information
SVM Stopped(ocumEvtVserver Down)	Risk	SVM	Warning
AV Server too Busy to Accept New Scan Request *	Risk	SVM	Error
No AV Server Connection for Virus Scan *	Incident	SVM	Critical
No AV Server Registered *	Risk	SVM	Error
No Responsive AV Server Connection *	Event	SVM	Information
Unauthorized User Attempt to AV Server *	Risk	SVM	Error
Virus Found By AV Server *	Risk	SVM	Error
SVM with Infinite Volume Storage Not Available(ocumEvtVserve rStorageNotAvailable)	Incident	SVMs with Infinite Volume	Critical
SVM with Infinite Volume Storage Partially Available(ocumEvtVserve rStoragePartiallyAvailable)	Risk	SVMs with Infinite Volume	Error

Event name(Trap name)	Impact level	Source type	Severity
SVM with Infinite Volume Namespace Mirror Constituents Having Availability Issues(ocumEvtVserverN sMirrorAvailabilityHavingI ssues)	Risk	SVMs with Infinite Volume	Warning

Impact area: capacity

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

Event name(Trap name)	Impact level	Source type	Severity
SVM with Infinite Volume Space Full(ocumEvtVserverFull)	Risk	SVM	Error
SVM with Infinite Volume Space Nearly Full(ocumEvtVserverNearl yFull)	Risk	SVM	Warning
SVM with Infinite Volume Snapshot Usage Limit Exceeded(ocumEvtVserv erSnapshotUsageExceed ed)	Risk	SVM	Warning
SVM with Infinite Volume Namespace Space Full(ocumEvtVserverNam espaceFull)	Risk	SVM	Error
SVM with Infinite Volume Namespace Space Nearly Full(ocumEvtVserverNam espaceNearlyFull)	Risk	SVM	Warning

Impact area: configuration

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

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

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
SVM IOPS Critical Threshold Breached(ocumSvmlopsI ncident)	Incident	SVM	Critical
SVM IOPS Warning Threshold Breached(ocumSvmlops Warning)	Risk	SVM	Warning
SVM MBps Critical Threshold Breached(ocumSvmMbps Incident)	Incident	SVM	Critical
SVM MBps Warning Threshold Breached(ocumSvmMbps Warning)	Risk	SVM	Warning
SVM Latency Critical Threshold Breached(ocumSvmLaten cylncident)	Incident	SVM	Critical
SVM Latency Warning Threshold Breached(ocumSvmLaten cyWarning)	Risk	SVM	Warning

SVM storage class events

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

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

Impact area: availability

Event name(Trap name)	Impact level	Source type	Severity
SVM Storage Class Not Available(ocumEvtVserve rStorageClassNotAvailabl e)	Incident	Storage class	Critical
SVM Storage Class Partially Available(ocumEvtVserve rStorageClassPartiallyAva ilable)	Risk	Storage class	Error

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
SVM Storage Class Space Nearly Full(ocumEvtVserverStora geClassNearlyFull)	Risk	Storage class	Warning
SVM Storage Class Space Full(ocumEvtVserverStora geClassFull)	Risk	Storage class	Error
SVM Storage Class Snapshot Usage Limit Exceeded(ocumEvtVserv erStorageClassSnapshot UsageExceeded)	Risk	Storage class	Warning

User and group quota events

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

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
User or Group Quota Disk Space Soft Limit Breached(ocumEvtUserO rGroupQuotaDiskSpaceS oftLimitBreached)	Risk	User or group quota	Warning
User or Group Quota Disk Space Hard Limit Reached(ocumEvtUserOr GroupQuotaDiskSpaceHa rdLimitReached)	Incident	User or group quota	Critical
User or Group Quota File Count Soft Limit Breached(ocumEvtUserO rGroupQuotaFileCountSof tLimitBreached)	Risk	User or group quota	Warning
User or Group Quota File Count Hard Limit Reached(ocumEvtUserOr GroupQuotaFileCountHar dLimitReached)	Incident	User or group quota	Critical

Volume events

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

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name(Trap name)	Impact level	Source type	Severity
Volume Restricted(ocumEvtVolum eRestricted)	Risk	Volume	Warning
Volume Offline(ocumEvtVolumeOf fline)	Incident	Volume	Critical
Volume Partially Available(ocumEvtVolume PartiallyAvailable)	Risk	Volume	Error

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

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
Thin-Provisioned Volume Space At Risk(ocumThinProvisionV olumeSpaceAtRisk)	Risk	Volume	Warning
Volume Space Full(ocumEvtVolumeFull)	Risk	Volume	Error
Volume Space Nearly Full(ocumEvtVolumeNearl yFull)	Risk	Volume	Warning

Event name(Trap name)	Impact level	Source type	Severity
Volume Logical Space Full *(volumeLogicalSpaceFull)	Risk	Volume	Error
Volume Logical Space Nearly Full *(volumeLogicalSpaceNe arlyFull)	Risk	Volume	Warning
Volume Logical Space Normal *(volumeLogicalSpaceAll OK)	Event	Volume	Information
Volume Snapshot Reserve Space Full(ocumEvtSnapshotFull)	Risk	Volume	Warning
Too Many Snapshot Copies(ocumEvtSnapshot TooMany)	Risk	Volume	Error
Volume Qtree Quota Overcommitted(ocumEvtV olumeQtreeQuotaOverco mmitted)	Risk	Volume	Error
Volume Qtree Quota Nearly Overcommitted(ocumEvtV olumeQtreeQuotaAlmost Overcommitted)	Risk	Volume	Warning
Volume Growth Rate Abnormal(ocumEvtVolum eGrowthRateAbnormal)	Risk	Volume	Warning
Volume Days Until Full(ocumEvtVolumeDays UntilFullSoon)	Risk	Volume	Error
Volume Space Guarantee Disabled(Not applicable)	Event	Volume	Information

Event name(Trap name)	Impact level	Source type	Severity
Volume Space Guarantee Enabled(Not Applicable)	Event	Volume	Information
Volume Space Guarantee Modified(Not applicable)	Event	Volume	Information
Volume Snapshot Reserve Days Until Full(ocumEvtVolumeSnap shotReserveDaysUntilFull Soon)	Risk	Volume	Error
FlexGroup Constituents Have Space Issues *(flexGroupConstituentsH aveSpaceIssues)	Risk	Volume	Error
FlexGroup Constituents Space Status All OK *(flexGroupConstituentsS paceStatusAllOK)	Event	Volume	Information
FlexGroup Constituents Have Inodes Issues *(flexGroupConstituentsH aveInodesIssues)	Risk	Volume	Error
FlexGroup Constituents Inodes Status All OK *(flexGroupConstituentsIn odesStatusAllOK)	Event	Volume	Information
WAFL Volume AutoSize Fail *	Risk	Volume	Error
WAFL Volume AutoSize Done *	Event	Volume	Information

Impact area: configuration

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

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

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
QoS Volume Max IOPS Warning Threshold Breached(ocumQosVolumeMaxlopsWarning)	Risk	Volume	Warning
QoS Volume Max MBps Warning Threshold Breached(ocumQosVolumeMaxMbpsWarning)	Risk	Volume	Warning
QoS Volume Max IOPS/TB Warning Threshold Breached(ocumQosVolumeMaxlopsPerTbWarning)	Risk	Volume	Warning
Volume IOPS Critical Threshold Breached(ocumVolumelo psIncident)	Incident	Volume	Critical
Volume IOPS Warning Threshold Breached(ocumVolumelo psWarning)	Risk	Volume	Warning
Volume MBps Critical Threshold Breached(ocumVolumeM bpsIncident)	Incident	Volume	Critical
Volume MBps Warning Threshold Breached(ocumVolumeM bpsWarning)	Risk	Volume	Warning
Volume Latency ms/op Critical Threshold Breached(ocumVolumeLa tencyIncident)	Incident	Volume	Critical

Event name(Trap name)	Impact level	Source type	Severity
Volume Latency ms/op Warning Threshold Breached(ocumVolumeLa tencyWarning)	Risk	Volume	Warning
Volume Cache Miss Ratio Critical Threshold Breached(ocumVolumeC acheMissRatioIncident)	Incident	Volume	Critical
Volume Cache Miss Ratio Warning Threshold Breached(ocumVolumeC acheMissRatioWarning)	Risk	Volume	Warning
Volume Latency and IOPS Critical Threshold Breached(ocumVolumeLa tencylopsIncident)	Incident	Volume	Critical
Volume Latency and IOPS Warning Threshold Breached(ocumVolumeLatencylopsWarning)	Risk	Volume	Warning
Volume Latency and MBps Critical Threshold Breached(ocumVolumeLa tencyMbpsIncident)	Incident	Volume	Critical
Volume Latency and MBps Warning Threshold Breached(ocumVolumeLa tencyMbpsWarning)	Risk	Volume	Warning
Volume Latency and Aggregate Perf. Capacity Used Critical Threshold Breached(ocumVolumeLa tencyAggregatePerfCapa cityUsedIncident)	Incident	Volume	Critical
Volume Latency and Aggregate Perf. Capacity Used Warning Threshold Breached(ocumVolumeLa tencyAggregatePerfCapa cityUsedWarning)	Risk	Volume	Warning

Event name(Trap name)	Impact level	Source type	Severity
Volume Latency and Aggregate Utilization Critical Threshold Breached(ocumVolumeLa tencyAggregateUtilizationI ncident)	Incident	Volume	Critical
Volume Latency and Aggregate Utilization Warning Threshold Breached(ocumVolumeLa tencyAggregateUtilization Warning)	Risk	Volume	Warning
Volume Latency and Node Perf. Capacity Used Critical Threshold Breached(ocumVolumeLa tencyNodePerfCapacityU sedIncident)	Incident	Volume	Critical
Volume Latency and Node Perf. Capacity Used Warning Threshold Breached(ocumVolumeLa tencyNodePerfCapacityU sedWarning)	Risk	Volume	Warning
Volume Latency and Node Perf. Capacity Used - Takeover Critical Threshold Breached(ocumVolumeLa tencyAggregatePerfCapa cityUsedTakeoverIncident)	Incident	Volume	Critical
Volume Latency and Node Perf. Capacity Used - Takeover Warning Threshold Breached(ocumVolumeLa tencyAggregatePerfCapa cityUsedTakeoverWarnin g)	Risk	Volume	Warning

Event name(Trap name)	Impact level	Source type	Severity
Volume Latency and Node Utilization Critical Threshold Breached(ocumVolumeLa tencyNodeUtilizationIncid ent)	Incident	Volume	Critical
Volume Latency and Node Utilization Warning Threshold Breached(ocumVolumeLa tencyNodeUtilizationWarn ing)	Risk	Volume	Warning

Volume move status events

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

Impact area: capacity

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

Description of event windows and dialog boxes

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

Event Retention Settings dialog box

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

You must have the OnCommand Administrator or Storage Administrator role.

Event Settings

You can configure the following options:

• Delete Information, Resolved, and Obsolete Events Older Than

Enables you to specify the retention period after which events that are marked as Information, Resolved, or Obsolete are removed from the management server.

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

· Delete Schedule

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

The default value is Daily.

Delete Now

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

Command buttons

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

Save and Close

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

Cancel

Cancels the recent changes and closes the dialog box.

Setup/Notifications page

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

You must have the OnCommand Administrator or Storage Administrator role.

Email

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

From Address

Specifies the email address from which the alert notification is sent. This value is also used as the from address for a report when shared. If the From Address is pre-filled with the address "OnCommand@localhost.com", you should change it to a real, working email address to make sure that all email notifications are delivered successfully.

SMTP Server

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

Host Name or IP Address

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

User Name

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

Password

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

Port

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

The default value is 25.

Use STARTTLS

Checking this box provides secure communication between the SMTP server and the management server by using the TLS/SSL protocols (also known as start tls and StartTLS).

· Use SSL

Checking this box provides secure communication between the SMTP server and the management server by using the SSL protocol.

SNMP

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

Version

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

Trap Destination Host

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

Outbound Trap Port

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

The default value is 162.

Community

The community string to access the host.

Engine ID

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

Username

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

Authentication Protocol

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

Authentication Password

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

Privacy Protocol

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

Privacy Password

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

Events inventory page

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

By default The information on this page is refreshed automatically every 5 minutes to ensure that the most current new events are displayed.

Filter components

Enable you to customize the information that is displayed in the events list. You can refine the list of events that are displayed using the following components:

• View menu to select from a pre-defined list of filter selections.

This includes items such as all active (new and acknowledged) events, active performance events, events assigned to me (the logged in user), and all events generated during all maintenance windows.

- Search pane to refine the list of events by entering full or partial terms.
- Filter button that launches the Filters pane so you can select from every available field and field attribute to refine the list of events.
- Time selector to refine the list of events by the time at which the event was triggered.

Command buttons

The command buttons enable you to perform the following tasks:

· Assign To

Enables you to select the user to whom the event is assigned. When you assign an event to a user, the user name and the time when you assigned the event is added in the events list for the selected events.

∘ Me

Assigns the event to the currently logged in user.

Another user

Displays the Assign Owner dialog box, which enables you to assign or reassign the event to other users. You can also unassign events by leaving the ownership field blank.

Acknowledge

Acknowledges the selected events.

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



You cannot acknowledge Information events.

Mark As Resolved

Enables you to change the event state to resolved.

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

Add Alert

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

Export

Enables you to export details of all events to a comma-separated values (.csv) file.

Column Selector

Enables you to choose the columns that display on the page and select the order in which they are displayed.

Events list

Displays details of all the events ordered by triggered time.

By default New and Acknowledged events for the previous seven days of severity type Critical, Error, and Warning are displayed.

Triggered Time

The time at which the event was generated.

Severity

The event severity: Critical (\mathbf{X}), Error ($\mathbf{0}$), Warning ($\mathbf{\Lambda}$), and Information ($\mathbf{0}$).

State

The event state: New, Acknowledged, Resolved, or Obsolete.

Impact Level

The event impact level: Incident, Risk, or Event.

Impact Area

The event impact area: Availability, Capacity, Performance, Protection, or Configuration.

Name

The event name.

You can select the event name to display the Event details page.

Source

The name of the object on which the event has occurred.

When a shared QoS policy breach occurs, only the workload object that is consuming the most IOPS or MBps is shown in this field. Additional workloads that are using this policy are displayed in the Event details page.

You can select the source name to display the health or performance details page for that object.

Source Type

The object type (for example, SVM, Volume, or Qtree) with which the event is associated.

Assigned To

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

Notes

The number of notes that are added for an event.

Days Outstanding

The number of days since the event was initially generated.

Assigned Time

The time that has elapsed since the event was assigned to a user. If the time elapsed exceeds a week, the timestamp when the event was assigned to a user is displayed.

Acknowledged By

The name of the user who acknowledged the event. The field is blank if the event is not acknowledged.

Acknowledged Time

The time that has elapsed since the event was acknowledged. If the time elapsed exceeds a week, the timestamp when the event was acknowledged is displayed.

Resolved By

The name of the user who resolved the event. The field is blank if the event is not resolved.

Resolved Time

The time that has elapsed since the event was resolved. If the time elapsed exceeds a week, the timestamp when the event was resolved is displayed.

Obsoleted Time

The time when the state of the event became Obsolete.

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 possible remediations to resolve the issue.

Event Name

The name of the event and the time the event was last seen.

For non-performance events, while the event is in the New or Acknowledged state the last seen information is not known and is therefore hidden.

Event Description

A brief description of the event.

In some cases a reason for the event being triggered is provided in the event description.

Component in Contention

For dynamic performance events, this section displays icons that represent the logical and physical components of the cluster. If a component is in contention, its icon is circled and highlighted red.

The following components may be displayed:

Network

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

Network Processing

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

QoS Policy

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

Cluster Interconnect

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

Data Processing

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

MetroCluster Resources

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

Aggregate or SSD Aggregate Ops

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

Cloud Latency

Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

Sync SnapMirror

Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

The Event Information, System Diagnosis, and Suggested Actions sections are described in other topics.

Command buttons

The command buttons enable you to perform the following tasks:

· Notes icon

Enables you to add or update a note about the event, and review all notes left by other users.

Actions menu

Assign to Me

Assigns the event to you.

Assign to Others

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

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

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

Acknowledge

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

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

Mark As Resolved

Enables you to change the event state to Resolved.

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

Add Alert

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

What the Event Information section displays

You use the Event Information section on the Event details page to view the details about a selected event, such as the event severity, impact level, impact area, and event source.

Fields that are not applicable to the event type are hidden. You can view the following event details:

Event Trigger Time

The time at which the event was generated.

State

The event state: New, Acknowledged, Resolved, or Obsolete.

Obsoleted Cause

The actions that caused the event to be obsoleted, for example, the issue was fixed.

Event Duration

For active (new and acknowledged) events, this is the time between detection and the time when the event was last analyzed. For obsolete events, this is the time between detection and when the event was resolved.

This field is displayed for all performance events, and for other event types only after they have been resolved or obsoleted.

· Last Seen

The date and time at which the event was last seen as active.

For performance events this value may be more recent than the Event Trigger Time as this field is updated after each new collection of performance data as long as the event is active. For other types of events, when in the New or Acknowledged state, this content is not updated and the field is therefore hidden.

Severity

The event severity: Critical (\bigotimes), Error (\bigcirc), Warning (\triangle), and Information (\bigcirc).

Impact Level

The event impact level: Incident, Risk, or Event.

Impact Area

The event impact area: Availability, Capacity, Performance, Protection, or Configuration.

Source

The name of the object on which the event has occurred.

When viewing the details for a shared QoS policy event, up to three of the workload objects that are consuming the most IOPS or MBps are listed in this field.

You can click the source name link to display the health or performance details page for that object.

Source Annotations

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

This field is displayed only for health events on clusters, SVMs, and volumes.

Source Groups

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

This field is displayed only for health events on clusters, SVMs, and volumes.

Source Type

The object type (for example, SVM, Volume, or Qtree) with which the event is associated.

On Cluster

The name of the cluster on which the event occurred.

You can click the cluster name link to display the health or performance details page for that cluster.

Affected Objects Count

The number of objects affected by the event.

You can click the object link to display the inventory page populated with the objects that are currently affected by this event.

This field is displayed only for performance events.

Affected Volumes

The number of volumes that are being affected by this event.

This field is displayed only for performance events on nodes or aggregates.

Triggered Policy

The name of the threshold policy that issued the event.

You can hover your cursor over the policy name to see the details of the threshold policy. For adaptive QoS policies the defined policy, block size, and allocation type (allocated space or used space) is also

displayed.

This field is displayed only for performance events.

Acknowledged by

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

· Resolved by

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

· Assigned to

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

Alert Settings

The following information about alerts is displayed:

· If there are no alerts associated with the selected event, an Add alert 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 Configuration/Alerting page by clicking the link to view more details about these alerts.

Alerts that are disabled are not displayed.

Last Notification Sent

The date and time at which the most recent alert notification was sent.

Sent Via

The mechanism that was used to send the alert notification: email or SNMP trap.

Previous Script Execution

The name of the script that was executed when the alert was generated.

What the System Diagnosis section displays

The System Diagnosis section of the Event details page provides information that can help you diagnose issues that may have been responsible for the event.

This area is displayed only for some events.

Some performance events provide charts that are relevant to the particular event that has been triggered.

Typically this includes and IOPS or MBps chart and a latency chart for the previous ten days. When arranged this way you can see which storage components are most affecting latency, or being affected by latency, when the event is active.

For dynamic performance events, the following charts are displayed:

- Workload Latency Displays the history of latency for the top victim, bully, or shark workloads at the component in contention.
- · Workload Activity Displays details about the workload usage of the cluster component in contention.
- Resource Activity Display historical performance statistics for the cluster component in contention.

Other charts are displayed when some cluster components are in contention.

Other events provide a brief description of the type of analysis the system is performing on the storage object. In some cases there will be one or more lines; one for each component that has been analyzed, for system-defined performance policies that analyze multiple performance counters. In this scenario, a green or red icon displays next to the diagnosis to indicate whether an issue was found, or not, in that particular diagnosis.

What the Suggested Actions section displays

The Suggested Actions section of the Event details page provides possible reasons for the event and suggests a few actions so that you can try to resolve the event on your own. The suggested actions are customized based on the type of event or type of threshold that has been breached.

This area is displayed only for some types of events.

In some cases there are **Help** links provided on the page that reference additional information for many suggested actions, including instructions for performing a specific action. Some of the actions may involve using Unified Manager, OnCommand System Manager, OnCommand Workflow Automation, ONTAP CLI commands, or a combination of these tools.

There are also some links provided in this help topic.

You should consider the actions suggested here as only a guidance in resolving this event. The action you take to resolve this event should be based on the context of your environment.

Configuration/Manage Events page

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

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

Command buttons

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

Disable

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

Enable

Enables selected events that you had chosen to disable previously.

Subscribe to EMS Events

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

Event Retention Settings

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

List view

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

Event

Displays the name of the event that is disabled.

Severity

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

Source Type

Displays the source type for which the event is generated.

Disable Events dialog box

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

You must have the OnCommand Administrator or Storage Administrator role.

Event Properties area

The Event Properties area specifies the following event properties:

Event Severity

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

Event Name Contains

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

· Matching events

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

Disable events

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

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

Command buttons

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

· Save and close

Disables the event type and closes the dialog box.

Cancel

Discards the changes and closes the dialog box.

Managing alerts

You can configure alerts to send notification automatically when specific events or events of certain severity types occur. You can also associate an alert with a script that is executed when an alert is triggered.

What alerts are

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

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

What information is contained in an alert email

Unified Manager alert emails provide the type of event, the severity of the event, the name of the policy that was breached to cause the event, and a description of the event. The email message also provides a hyperlink for each event that enables you to view the details page for the event in the UI.

Alert emails are sent to all users who have subscribed to receive alerts.

If a performance counter or capacity value has a large change during a collection period, it could cause both a critical and a warning event to be triggered at the same time for the same threshold policy. In this case, you may receive one email for the warning event and one for the critical event. This is because Unified Manager enables you to subscribe separately to receive alerts for warning and critical threshold breaches.



After upgrading to Unified Manager 7.2, or greater, links to events and alerts from emails that were send from older versions of Unified Manager will no longer work because of a change in the event and alert URLs.

A sample alert email is shown below:

From: 10.11.12.13@company.com Sent: Tuesday, May 1, 2018 7:45 PM

To: sclaus@company.com; user1@company.com

Subject: Alert from OnCommand Unified Manager: Thin-Provisioned Volume Space At Risk (State: New)

A risk was generated by 10.11.12.13 that requires your attention.

Risk - Thin-Provisioned Volume Space At Risk

Impact Area - Capacity Severity - Warning State - New

Source - svm_n1:/sm_vol_23 Cluster Name - fas3250-39-33-37

Cluster FQDN - fas3250-39-33-37-cm.company.com

Trigger Condition - The thinly provisioned capacity of the volume is 45.73% of the available space on the host aggregate. The capacity of the volume is at risk because of aggregate capacity issues.

Event details:

https://10.11.12.13:443/events/94

Source details:

https://10.11.12.13:443/health/volumes/106

Alert details:

https://10.11.12.13:443/alerting/1

Adding alerts

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

Before you begin

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

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

About this task

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

Steps

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

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

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

- 5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.
 - To select more than one event, press the Ctrl key while you make your selections.
- 6. Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.



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

You can also choose to notify users through SNMP traps.

7. Click Save.

Example of adding an alert

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

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

Perform the following steps in the Add Alert dialog box:

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

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

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

Guidelines for adding alerts

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

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

Alert description

You should provide a description for the alert so that it helps you track your alerts effectively.

Resources

You should decide which physical or logical resource requires an alert. You can include and exclude resources, as required. For example, if you want to closely monitor your aggregates by configuring an alert, you must select the required aggregates from the list of resources.

Event severity

You should decide if an event of a specified severity type (Critical, Error, Warning) should trigger the alert and, if so, which severity type.

· Event name

If you add an alert based on the type of event generated, you should decide which events require an alert

Actions

You must provide the user names and email addresses of the users who receive the notification. You can also specify an SNMP trap as a mode of notification. You can associate your scripts to an alert so that they

are executed when an alert is generated.

Notification frequency

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

· Execute Script

You can associate your script with an alert. Your script is executed when the alert is generated.

Adding alerts for performance events

You can configure alerts for individual performance events just like any other events received by Unified Manager. Additionally, if you want to treat all performance events alike and have email sent to the same person, you can create a single alert to notify you when any critical or warning performance events are triggered.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The example below shows how to create an event for all critical latency, IOPS, and MBps events. You can use this same methodology to select events from all performance counters, and for all warning events.

Steps

- 1. In the left navigation pane, click Configuration > Alerting.
- 2. In the Configuration/Alerting page, click Add.
- 3. In the Add Alert dialog box, click Name, and enter a name and description for the alert.
- 4. Do not select any resources on the **Resources** page.

Because no resources are selected, the alert is applied to all clusters, aggregates, volumes, and so on, for which these events are received.

- 5. Click **Events** and perform the following actions:
 - a. In the Event Severity list, select Critical.
 - b. In the Event Name Contains field, enter latency and then click the arrow to select all the matching events.
 - c. In the Event Name Contains field, enter iops and then click the arrow to select all the matching events.
 - d. In the Event Name Contains field, enter mbps and then click the arrow to select all the matching events.
- 6. Click **Actions** and then select the name of the user who will receive the alert email in the **Alert these** users field.

- 7. Configure any other options on this page for issuing SNMP taps and executing a script.
- 8. Click Save.

Excluding disaster recovery destination volumes from generating alerts

When configuring volume alerts you can specify a string in the Alert dialog box that identifies a volume or group of volumes. If you have configured disaster recovery for SVMs, however, the source and destination volumes have the same name, so you will receive alerts for both volumes.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can disable alerts for disaster recovery destination volumes by excluding volumes that have the name of the destination SVM. This is possible because the identifier for volume events contains both the SVM name and volume name in the format "<svm_name>:/<volume_name>".

The example below shows how to create alerts for volume "vol1" on the primary SVM"`vs1`", but exclude the alert from being generated on a volume with the same name on SVM"`vs1-dr`".

Perform the following steps in the Add Alert dialog box:

Steps

- 1. Click **Name** and enter a name and description for the alert.
- 2. Click **Resources**, and then select the **Include** tab.
 - a. Select **Volume** from the drop-down list, and then enter vol1 in the **Name contains** field to display the volumes whose name contains "vol1".
 - b. Select [All Volumes whose name contains 'vol1'] from the Available Resources area, and move it to the Selected Resources area.
- 3. Select the **Exclude** tab, select **Volume**, enter vs1-dr in the **Name contains** field, and then click **Add**.

This excludes the alert from being generated for volume "vol1" on SVM"'vs1-dr'".

- 4. Click **Events** and select the event or events that you want to apply to the volume or volumes.
- 5. Click **Actions** and then select the name of the user who will receive the alert email in the **Alert these** users field.
- 6. Configure any other options on this page for issuing SNMP traps and executing a script, and then click **Save**.

Testing alerts

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

Before you begin

 You must have configured notification settings such as the email address of the recipients, SMTP server, and SNMP trap.

The Unified Manager server can use these settings to send notifications to users when an event is generated.

- You must have assigned a script and configured the script to run when the alert is generated.
- You must have the OnCommand Administrator role.

Steps

- 1. In the left navigation pane, click **Configuration > Alerting**.
- 2. In the Configuration/Alerting page, select the alert that you want to test, and then click Test.

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

Viewing alerts

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

Before you begin

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

Steps

1. In the left navigation pane, click **Configuration > Alerting**.

The list of alerts is displayed in the Configuration/Alerting page.

Editing alerts

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

Before you begin

You must have the OnCommand Administrator role.

Steps

- 1. In the left navigation pane, click **Configuration > Alerting**.
- 2. In the Configuration/Alerting page, select the alert that you want to edit, and click Edit.
- 3. In the Edit Alert dialog box, edit the name, resources, events, and actions sections, as required.

You can change or remove the script that is associated with the alert.

4. Click Save.

Deleting alerts

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

Before you begin

You must have the OnCommand Administrator role.

Steps

- 1. In the left navigation pane, click **Configuration > Alerting**.
- 2. On the Configuration/Alerting page, select the alerts that you want to delete, and click Delete.
- Click Yes to confirm the delete request.

Description of alert windows and dialog boxes

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

Configuration/Alerting page

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

You must have the OnCommand Administrator or Storage Administrator role.

Command buttons

Add

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

• Edit

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

Delete

Deletes the selected alerts.

Enable

Enables the selected alerts to send notifications.

Disable

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

Test

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

List view

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

Status

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

Alert

Displays the name of the alert.

Description

Displays a description for the alert.

Notification Method

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

Notification Frequency

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

Details area

The details area provides more information about the selected alert.

Alert Name

Displays the name of the alert.

Alert Description

Displays a description for the alert.

Events

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

Resources

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

Includes

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

Excludes

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

Notification Method

Displays the notification method for the alert.

Notification Frequency

Displays the frequency with which the management server continues to send alert notifications until the event is acknowledged, resolved, or moved to the Obsolete state.

Script Name

Displays the name of the script associated with the selected alert. This script is executed when an alert is generated.

Email Recipients

Displays the email addresses of users who receive the alert notification.

Add Alert dialog box

You can create alerts to notify you when a particular event is generated, so that you can address the issue quickly and thereby minimize impact to your environment. You can create alerts for a single resource or a set of resources, and for events of a particular severity type. You can also specify the notification method and frequency of the alerts.

You must have the OnCommand Administrator or Storage Administrator role.

Name

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

Alert Name

Enables you to specify an alert name.

Alert Description

Enables you to specify a description for the alert.

Resources

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

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

Include

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

Exclude

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

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

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

Events

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

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

Event Severity

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

Event Name Contains

Enables you to select events whose name contains specified characters.

Actions

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

Alert these users

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

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

Notification Frequency

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

You can choose the following notification methods:

Notify only once

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

Issue SNMP trap

Selecting this box enables you to specify whether SNMP traps should be sent to the globally configured SNMP host.

Execute Script

Enables you to add your custom script to the alert. This script is executed when an alert is generated.

Command buttons

Save

Creates an alert and closes the dialog box.

Cancel

Discards the changes and closes the dialog box.

Edit Alert dialog box

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

Name

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

Alert Name

Enables you to edit the alert name.

Alert Description

Enables you to specify a description for the alert.

Alert State

Enables you to enable or disable the alert.

Resources

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

Include

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

Exclude

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



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

Events

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

Event Severity

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

Event Name Contains

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

Actions

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

Alert these users

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

Notification Frequency

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

You can choose the following notification methods:

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

Issue SNMP trap

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

Execute Script

Enables you to associate a script with the alert. This script is executed when an alert is generated.

Command buttons

Save

Saves the changes and closes the dialog box.

Cancel

Discards the changes and closes the dialog box.

Managing scripts

You can use scripts to automatically modify or update multiple storage objects in Unified Manager. The script is associated with an alert. When an event triggers an alert, the script is executed. You can upload custom scripts and test their execution when an alert is generated.

How scripts work with alerts

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

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

The script uses the following arguments for execution:

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

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

Example for obtaining arguments from scripts

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

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

-eventID : 290

-eventSourceID: 4138

Adding scripts

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

Before you begin

- You must have created and saved the scripts that you want to add to the Unified Manager server.
- The supported file formats for scripts are Perl, Shell, PowerShell, and .bat files.
 - For Perl scripts, Perl must be installed on the Unified Manager server. If Perl was installed after Unified Manager, you must restart the Unified Manager server.
 - For PowerShell scripts, the appropriate PowerShell execution policy must be set on the server so that the scripts can be executed.



If your script creates log files to track the alert script progress, you must make sure that the log files are not created anywhere within the Unified Manager installation folder.

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

About this task

You can upload custom scripts and gather event details about the alert.

Steps

- 1. In the toolbar, click , and then click **Scripts** in the left Management menu.
- 2. In the Management/Scripts page, click Add.
- 3. In the **Add Script** dialog box, click **Browse** to select your script file.
- 4. Enter a description for the script that you select.
- 5. Click Add.

Deleting scripts

You can delete a script from Unified Manager when the script is no longer required or valid.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The script must not be associated with an alert.

Steps

- 1. In the toolbar, click , and then click **Scripts** in the left Management menu.
- 2. In the Management/Scripts page, select the script that you want to delete, and then click Delete.
- 3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Testing script execution

You can verify that your script is executed correctly when an alert is generated for a storage object.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have uploaded a script in the supported file format to Unified Manager.

Steps

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

То	Do this
Add an alert	a. In the Configuration/Alerting page, click Add .
	b. In the Actions section, associate the alert with your test script.
Edit an alert	a. In the Configuration/Alerting page, select an alert, and then click Edit .
	b. In the Actions section, associate the alert with your test script.

- 4. Click Save.
- 5. In the Configuration/Alerting page, select the alert that you added or modified, and then click Test.

The script is executed with the "-test" argument, and a notification alert is sent to the email addresses that were specified when the alert was created.

Description of script windows and dialog boxes

The Management/Scripts page enables you to add scripts to Unified Manager.

Management/Scripts page

The Management/Scripts page enables you to add your custom scripts to Unified Manager. You can associate these scripts with alerts to enable automatic reconfiguration

of storage objects.

The Management/Scripts page enables you to add or delete scripts from Unified Manager.

Command buttons

Add

Displays the Add Script dialog box, which enables you to add scripts.

Delete

Deletes the selected script.

List view

The list view displays, in tabular format, the scripts that you added to Unified Manager.

Name

Displays the name of the script.

Description

Displays the description of the script.

Add Script dialog box

The Add Script dialog box enables you to add scripts to Unified Manager. You can configure alerts with your scripts to automatically resolve events that are generated for storage objects.

You must have the OnCommand Administrator or Storage Administrator role.

Select Script File

Enables you to select a script for the alert.

Description

Enables you to specify a description for the script.

Supported Unified Manager CLI commands

As a storage administrator you can use the CLI commands to perform queries on the storage objects; for example, on clusters, aggregates, volumes, qtrees, and LUNs. You can use the CLI commands to query the Unified Manager internal database and the ONTAP database. You can also use CLI commands in scripts that are executed at the beginning or end of an operation or are executed when an alert is triggered.

All commands must be preceded with the command um cli login and a valid user name and password for authentication.

CLI command	Description	Output
<pre>um run cmd [-t <timeout>] <cluster> <command/></cluster></timeout></pre>	The simplest way to run a command on one or more hosts. Mainly used for alert scripting to get or perform an operation on ONTAP. The optional timeout argument sets a maximum time limit (in seconds) for the command to complete on the client. The default is 0 (wait forever).	As received from ONTAP.
um run query <sql command=""></sql>	Executes an SQL query. Only queries that read from the database are allowed. Any update, insert, or delete operations are not supported.	Results are displayed in a tabular form. If an empty set is returned, or if there is any syntax error or bad request, it displays the appropriate error message.
<pre>um datasource add -u <username> -P <password> [-t <protocol>] [-p <port>] <hostname-or-ip></hostname-or-ip></port></protocol></password></username></pre>	Adds a datasource to the list of managed storage systems. A datasource describes how connections to storage systems are made. The options -u (username) and -P (password) must be specified when adding a datasource. The option -t (protocol) specifies the protocol used to communicate with the cluster (http or https). If the protocol is not specified, then both protocols will be attempted The option -p (port) specifies the port used to communicate with the cluster. If the port is not specified, then the default value of the appropriate protocol will be attempted. This command can be executed only by the storage admin.	Prompts for the user accept the certificate and prints the corresponding message.
um datasource list [<datasource-id>]</datasource-id>	Displays the datasources for managed storage systems.	Displays the following values in tabular format: ID Address Port, Protocol Acquisition Status, Analysis Status, Communication status, Acquisition Message, and Analysis Message.

CLI command	Description	Output
<pre>um datasource modify [-h <hostname-or-ip>] [-u <username>] [-P <password>] [-t <protocol>] [-p <port>] <datasource-id></datasource-id></port></protocol></password></username></hostname-or-ip></pre>	Modifies one or more datasource options. Can be executed only by the storage admin.	Displays the corresponding message.
um datasource remove <datasource-id></datasource-id>	Removes the datasource from Unified Manager.	Displays the corresponding message.
<pre>um option list [<option>]</option></pre>	Lists options.	Displays the following values in tabular format: Name, Value, Default Value, and Requires Restart.
<pre>um option set <option- name="">=<option-value> [<option-name>=<option- value="">]</option-></option-name></option-value></option-></pre>	Sets one or more options. The command can be executed only by the storage admin.	Displays the corresponding message.
um version	Displays the Unified Manager software version .	Version ("7.0")
<pre>um lun list [-q] [-ObjectType <object-id>]</object-id></pre>	Lists the LUNs after filtering on the specified objectq is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: um lun list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the LUNs within the cluster with ID 1.	Displays the following values in tabular format: ID and LUN path.
<pre>um svm list [-q] [-ObjectType <object-id>]</object-id></pre>	Lists the SVMs after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: um svm list-cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the SVMs within the cluster with ID 1.	Displays the following values in tabular format: Name and Cluster ID.

CLI command	Description	Output
<pre>um qtree list [-q] [-ObjectType <object-id>]</object-id></pre>	Lists the qtrees after filtering on the specified objectq is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: um qtree list-cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the qtrees within the cluster with ID 1.	Displays the following values in tabular format: Qtree ID and Qtree Name.
<pre>um disk list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the disks after filtering on the specified object. ObjectType can be disk, aggr, node, cluster. For example: um disk list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the disks within the cluster with ID 1.	Displays the following values in tabular format ObjectType and object-id.
<pre>um cluster list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the clusters after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, lun, qtree, volume, quota, svm. For example:um cluster list -aggr 1 In this example, "-aggr" is the objectType and "1" is the objectId. The command lists the cluster to which the aggregate with ID 1 belongs.	Displays the following values in tabular format: Name, Full Name, Serial Number, Datasource Id, Last Refresh Time, and Resource Key.
<pre>um cluster node list [-q] [-ObjectType <object-id>]</object-id></pre>	Lists the cluster nodes after filtering on the specified object. ObjectType can be disk, aggr, node, cluster. For example: um cluster node list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the nodes within the cluster with ID 1.	Displays the following values in tabular format Name and Cluster ID.

CLI command	Description	Output
<pre>um volume list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the volumes after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, svm, aggregate. For example: um volume list-cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the volumes within the cluster with ID 1.	Displays the following values in tabular format Volume ID and Volume Name.
<pre>um quota user list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the quota users after filtering on the specified object. ObjectType can be qtree, cluster, volume, quota, svm. For example: um quota user list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the quota users within the cluster with ID 1.	Displays the following values in tabular format ID, Name, SID and Email.
<pre>um aggr list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the aggregates after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, volume. For example: um aggr list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the aggregates within the cluster with ID 1.	Displays the following values in tabular format Aggr ID, and Aggr Name.
um event ack <event-ids></event-ids>	Acknowledges one or more events.	Displays the corresponding message.
<pre>um event resolve <event- ids=""></event-></pre>	Resolves one or more events.	Displays the corresponding message.
um event assign -u <username> <event-id></event-id></username>	Assigns an event to a user.	Displays the corresponding message.
<pre>um event list [-s <source/>] [-S <event- state-filter-list="">] [<event-id>]</event-id></event-></pre>	Lists the events generated by the system or user. Filters events based on source, state, and IDs.	Displays the following values in tabular format Source, Source type, Name, Severity, State, User and Timestamp.

CLI command	Description	Output
um cli login -u <username> [-p <password></password></username>	Logs in to the CLI. The session expires after three hours from the time of login, after which the user must login again.	Displays the corresponding message.
um cli logout	Logs out of the CLI.	Displays the corresponding message.
<pre>um backup restore -f</pre>	Restores a database backup using .7z files.	Displays the corresponding message.
um help	Displays all first level subcommands.	Displays all first level subcommands.

Managing health thresholds

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

What storage capacity health thresholds are

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

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

Configuring global health threshold settings

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

About this task

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

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

Choices

· Configuring global aggregate health threshold values

You can configure the health threshold settings for capacity, growth, and Snapshot copies for all aggregates to track any threshold breach.

· Configuring global volume health threshold values

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

Configuring global qtree health threshold values

You can edit the health threshold settings for capacity for all gtrees to track any threshold breach.

Editing lag health threshold settings for unmanaged protection relationships

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

Configuring global aggregate health threshold values

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

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

When Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.



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

Steps

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

Configuring global volume health threshold values

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

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

Steps

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

Configuring global qtree health threshold values

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Events are generated for a qtree only when a Qtree quota or a Default quota has been set on the qtree. Events are not generated if the space defined in a User quota or Group quota has exceeded the threshold.

Steps

- 1. In the left navigation pane, click **Configuration > Health Thresholds**.
- 2. In the Configuration/Health Thresholds page, click Qtrees.

- 3. Configure the appropriate capacity threshold values.
- 4. Click Save.

Editing lag health threshold settings for unmanaged protection relationships

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The lag time must be no more than the defined transfer schedule interval. For example, if the transfer schedule is hourly, then the lag time must not be more than one hour. The lag threshold specifies a percentage that the lag time must not exceed. Using the example of one hour, if the lag threshold is defined as 150%, then you will receive an event when the lag time is more than 1.5 hours.

The settings described in this task are applied globally to all unmanaged protection relationships. The settings cannot be specified and applied exclusively to one unmanaged protection relationship.

Steps

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

Editing individual aggregate health threshold settings

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

When Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Steps

- 1. In the left navigation pane, click **Health > Aggregates**.
- 2. In the Health/Aggregates inventory page, select one or more aggregates and then click Edit Thresholds.
- 3. In the **Edit Aggregate Thresholds** dialog box, edit the threshold settings of one of the following: capacity, growth, or Snapshot copies by selecting the appropriate check box and then modifying the settings.
- 4. Click Save.

Editing individual volume health threshold settings

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

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

Steps

- 1. In the left navigation pane, click **Health > Volumes**.
- 2. In the **Health/Volumes** inventory page, select one or more volumes and then click **Edit Thresholds**.
- 3. In the **Edit Volume Thresholds** dialog box, edit the threshold settings of one of the following: capacity, Snapshot copies, qtree quota, growth, or inodes by selecting the appropriate check box and then modifying the settings.
- 4. Click Save.

Editing individual qtree health threshold settings

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the **Health/Storage Virtual Machines** inventory page, select the SVM on which the gtree resides.
- In the Health/Storage Virtual Machine details page, click the Qtrees tab.
- 4. Select one or more gtrees and then click **Edit Thresholds**.
- 5. In the **Edit Qtree Thresholds** dialog box, change the capacity thresholds for the selected qtree or qtrees and click **Save**.

Description of health thresholds pages

You can use the appropriate Configuration/Health Thresholds page to configure global health threshold values for aggregates and volumes, and configure global lag warning and error threshold values for unmanaged protection relationships.

Configuration/Health Thresholds page for Aggregates

The Configuration/Health Thresholds page for aggregates enables you to configure global health threshold values for monitored aggregates. When you configure the options globally, the default values of all objects are modified. However, if the default values have been changed at the object level, the global values are not modified.

You must have the OnCommand Administrator or Storage Administrator role.

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

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

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

Capacity area

The Capacity area enables you to set the following aggregate capacity threshold conditions. Note that when Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Space Nearly Full

Specifies the percentage at which an aggregate is considered to be nearly full:

· Default value: 80 percent

The value for this threshold must be lower than the value for the Aggregate Full threshold for the management server to generate an event.

Event generated: Aggregate Nearly Full

Event severity: Warning

Space Full

Specifies the percentage at which an aggregate is considered full:

· Default value: 90 percent

· Event generated: Aggregate Full

Event severity: Error

Nearly Overcommitted

Specifies the percentage at which an aggregate is considered to be nearly overcommitted:

Default value: 95 percent

The value for this threshold must be lower than the value for the Aggregate Overcommitted Full threshold for the management server to generate an event.

Event generated: Aggregate Nearly Overcommitted

· Event severity: Warning

Overcommitted

Specifies the percentage at which an aggregate is considered to be overcommitted:

Default value: 100 percent

Event generated: Aggregate Overcommitted

· Event severity: Error

· Days Until Full

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

Default value: 7

· Event generated: Aggregate Days Until Full

Event severity: Error

Growth area

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

Growth Rate

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

· Default value: 1 percent

Event generated: Aggregate Growth Rate Abnormal

· Event severity: Warning

Growth Rate Sensitivity

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

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

Default value: 2



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

Snapshot copies area

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

Snapshot Reserve Full

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

Default value: 90 percent

Event generated: Aggregate Snapshot Reserve Full

Event severity: Warning

Command buttons

Restore to Factory Defaults

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

Save

Saves the configuration settings for the selected option.

Configuration/Health Thresholds page for Volumes

The Configuration/Health Thresholds page for Volumes enables you to configure global health threshold values for monitored volumes. You can set thresholds for individual volumes or for all the volumes globally. When you configure the options globally, the default values of all objects are modified. However, if the default values have been changed at the object level, the global values are not modified.

You must have the OnCommand Administrator or Storage Administrator role.

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

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

Capacity area

The Capacity area enables you to set the following volume capacity threshold conditions. Note that when Autogrow is enabled on a volume that capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Space Nearly Full

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

Default value: 80 percent

The value for this threshold must be lower than the value for the Volume Full threshold in order for the management server to generate an event.

Event generated: Volume Nearly Full

Event severity: Warning

Space Full

Specifies the percentage at which a volume is considered full:

· Default value: 90 percent

Event generated: Volume Full

Event Severity: Error

Days Until Full

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

Default value: 7

Event generated: Volume Days Until Full

· Event severity: Error

Snapshot copies area

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

Snapshot Reserve Full

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

· Default value: 90 percent

Event generated: Volume Snapshot Reserve Full

· Event severity: Error

Days Until Full

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

Default value: 7

Event generated: Volume Snapshot Reserve Days Until Full

· Event severity: Error

Count

Specifies the number of Snapshot copies on a volume that are considered to be too many:

Default value: 250

Event generated: Too Many Snapshot Copies

Event severity: Error

Qtree Quota area

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

Nearly Overcommitted

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

Default value: 95 percent

Event generated: Volume Qtree Quota Nearly Overcommitted

Event severity: Warning

Overcommitted

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

Default value: 100 percent

Event generated: Volume Qtree Quota Overcommitted

· Event severity: Error

Growth area

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

· Growth Rate

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

· Default value: 1 percent

Event generated: Volume Growth Rate Abnormal

Event severity: Warning

Growth Rate Sensitivity

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

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

Default value: 2



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

Inodes area

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

Nearly Full

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

· Default value: 80 percent

Event generated: Inodes Nearly Full

Event severity: Warning

Full

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

Default value: 90 percent

Event generated: Inodes Full

Event severity: Error

Command buttons

Restore to Factory Defaults

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

Save

Saves the configuration settings for the selected option.

Lag Thresholds for Unmanaged Relationships page

The Lag Thresholds for Unmanaged Relationships page enables you to configure global lag warning and error threshold values for unmanaged protection relationships so that you are notified and can take action when lag or threshold errors occur. Changes to these settings are applied during the next scheduled update.

You must have the OnCommand Administrator or Storage Administrator role.

Events are generated when a threshold is breached. You can take corrective actions for such events. Lag threshold settings for unmanaged relationships are enabled by default.

The lag threshold specifies a percentage that the lag time must not exceed. Using an example of one hour, if the lag threshold is defined as 150%, then you will receive an event when the lag time is more than 1.5 hours.

Lag Thresholds for Unmanaged Relationships area

The Lag area enables you set unmanaged relationship lag thresholds for the following conditions:

Warning

Specifies the percentage at which the lag duration equals or exceeds the lag warning threshold:

• Default value: 150 percent

Events generated: SnapMirror Relationship Lag Warning or SnapVault Relationship Lag Warning

Event severity: Warning

Error

Specifies the percentage at which the lag duration equals or exceeds the lag error threshold:

• Default value: 250 percent

Events generated: SnapMirror Relationship Lag Error or SnapVault Relationship Lag Error

Event severity: Error

Command buttons

Restore to Factory Defaults

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

Save

Saves the configuration settings for the selected option.

Configuration/Health Thresholds page for Qtrees

The Configuration/Health Thresholds page for Qtrees enables you to configure global health threshold values for monitored qtrees. Events are generated for a qtree only when a Qtree quota or a Default quota has been set on the qtree. Events are not generated if the space defined in a User quota or Group quota has exceeded the threshold.

You must have the OnCommand Administrator or Storage Administrator role.

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

Capacity area

The Capacity area enables you to set the following qtree capacity threshold conditions.

Space Nearly Full

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

· Default value: 80 percent

The value for this threshold must be lower than the value for the Qtree Full threshold.

Event generated: Qtree Nearly Full

Event severity: Warning

Space Full

Specifies the percentage at which a qtree is considered full:

· Default value: 90 percent

Event generated: Qtree Full

Event severity: Error

Command buttons

Restore to Factory Defaults

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

Save

Saves the configuration settings for the selected option.

Edit Aggregate Thresholds dialog box

You can configure alerts to send notifications when an event related to an aggregate's capacity is generated, and you can take corrective actions for the event. For example, for the Aggregate Full threshold, you can configure an alert to send notification when the condition persists over a specified period.

You must have the OnCommand Administrator or Storage Administrator role.

The Edit Aggregate Thresholds dialog box enables you to configure aggregate-level thresholds that are applied to selected aggregates. If you configure aggregate-level thresholds, they take priority over the global-level threshold values. You can configure threshold settings for capacity, growth, and Snapshot copies at the aggregate level. If these settings are not configured, the global threshold values are applied.



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

Capacity area

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

Space Nearly Full

Specifies the percentage at which an aggregate is considered to be nearly full. It also displays the size of the aggregate corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which an aggregate is considered full. It also displays the size of the aggregate corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Nearly Overcommitted

Specifies the percentage at which an aggregate is considered to be nearly overcommitted.

Overcommitted

Specifies the percentage at which an aggregate is considered to be overcommitted.

Days Until Full

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

Growth area

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

Growth Rate

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

Growth Rate Sensitivity

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

A lower value for growth rate sensitivity indicates that the aggregate is highly sensitive to changes in the growth rate.



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

Snapshot copies area

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

· Snapshot Reserve Full

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

You can also use the slider to set the threshold value.

Command buttons

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

Restore to Defaults

Enables you to restore the aggregate-level threshold values to the global values.

Save

Saves all the threshold settings.

· Save and Close

Saves all the threshold settings and then closes the dialog box.

Cancel

Ignores the changes (if any) to the threshold settings and closes the dialog box.

Edit Volume Thresholds dialog box

You can configure alerts to send notifications when an event related to a volume's capacity is generated, and you can take corrective actions for the event. For example, for the Volume Full threshold, you can configure an alert to send notification when the condition persists over a specified period.

You must have the OnCommand Administrator or Storage Administrator role.

The Edit Volume Thresholds dialog box enables you to configure volume-level thresholds that are applied to the selected volumes. When thresholds are configured at the volume level, they take priority over the group-level thresholds or the global-level threshold values.

You can configure threshold settings for capacity, Snapshot copies, qtree quota, growth, and inodes at the volume level. When a group action of volume threshold type is configured, the group action threshold values are used for settings that are not configured at the volume level. When no group action of volume threshold type is configured, areas in Edit Volume Thresholds dialog box that are not configured, use global threshold values.

Capacity area

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

Space Nearly Full

Specifies the percentage at which a volume is considered to be nearly full. It also displays the size of the volume corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which a volume is considered full. It also displays the size of the volume corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

· Days Until Full

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

Snapshot Copies

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

Snapshot Reserve Full

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

· Days Until Full

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

Count

Specifies the number of Snapshot copies on a volume that are considered to be too many.

Qtree Quota area

The Qtree Quota area enables you to set the following qtree quota threshold conditions for the selected volumes:

Nearly Overcommitted

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

Overcommitted

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

Growth area

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

Growth Rate

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

Growth Rate Sensitivity

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

A lower value for growth rate sensitivity indicates that the volume is highly sensitive to changes in the growth rate.



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

Inodes area

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

Nearly Full

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

You can also use the sliders to set the threshold value.

• Full

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

You can also use the sliders to set the threshold value.

Command buttons

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

· Restore to Defaults

Enables you to restore the threshold values to one of the following:

- Group values, if the volume belongs to a group and that group has a volume threshold action type.
- Global values, if the volume does not belong to any group or if it belongs to a group that does not have a volume threshold action type.

Save

Saves all the threshold settings.

· Save and Close

Saves all the threshold settings and then closes the dialog box.

Cancel

Ignores the changes (if any) to the threshold settings and closes the dialog box.

Edit Qtree Thresholds dialog box

You can configure alerts to send notifications when an event related to a qtree's capacity is generated, and you can take corrective actions for the event. For example, for the Qtree Full threshold, you can configure an alert to send notification when the condition persists over a specified period.

You must have the OnCommand Administrator or Storage Administrator role.

The Edit Qtree Thresholds dialog box enables you to configure qtree-level thresholds that are applied to the selected qtrees. When thresholds are configured at the qtree level, they take priority over the group-level thresholds or the global-level threshold values.

You can configure threshold settings for capacity at the qtree level. When a group action of qtree threshold type is configured, the group action threshold values are used for settings that are not configured at the qtree level. When no group action of qtree threshold type is configured, areas in Edit Qtree Thresholds dialog box that are not configured, use global threshold values.

Capacity area

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

Space Nearly Full

Specifies the percentage at which a qtree is considered to be nearly full. It also displays the size of the qtree corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which a qtree is considered full. It also displays the size of the qtree corresponding to the specified threshold value.

You can also use the slider to set the threshold value.

Command buttons

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

· Restore to Defaults

Enables you to restore the threshold values to one of the following:

- Group values, if the gtree belongs to a group and that group has a gtree threshold action type.
- Global values, if the qtree does not belong to any group or if it belongs to a group that does not have a
 qtree threshold action type.

Save

Saves all the threshold settings.

· Save and Close

Saves all the threshold settings and then closes the dialog box.

Cancel

Ignores the changes (if any) to the threshold settings and closes the dialog box.

Managing user-defined performance thresholds

Performance threshold policies enable you to determine the point at which Unified Manager generates an event to inform system administrators about issues that could be impacting workload performance. These threshold policies are known as *user-defined* performance thresholds.

This release supports user-defined, system-defined, and dynamic performance thresholds. With dynamic and system-defined performance thresholds, Unified Manager analyzes the workload activity to determine the appropriate threshold value. With user-defined thresholds, you can define the upper performance limits for many performance counters and for many storage objects.



System-defined performance thresholds and dynamic performance thresholds are set by Unified Manager and are not configurable. If you are receiving unnecessary events from any system-defined performance threshold policies, you can disable individual policies from the Configuration/Manage Events page.

How user-defined performance threshold policies work

You set performance threshold policies on storage objects (for example, on aggregates and volumes) so that an event can be sent to the storage administrator to inform the administrator that the cluster is experiencing a performance issue.

You create a performance threshold policy for a storage object by:

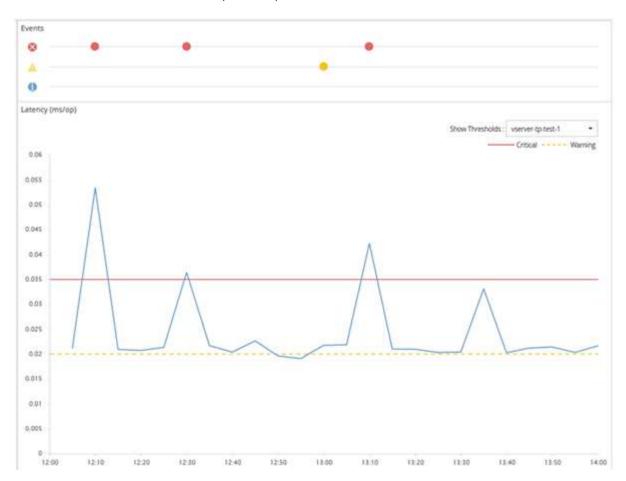
- · Selecting a storage object
- · Selecting a performance counter associated with that object
- Specifying values that define the performance counter upper limits that are considered warning and critical situations
- · Specifying a time period that defines how long the counter must exceed the upper limit

For example, you can set a performance threshold policy on a volume so that you receive a critical event notification whenever IOPS for that volume exceeds 750 operations per second for 10 consecutive minutes. This same threshold policy can also specify that a warning event be sent when IOPS exceeds 500 operations per second for 10 minutes.



The current release provides thresholds that send events when a counter value exceeds the threshold setting. You cannot set thresholds that send events when a counter value falls below a threshold setting.

An example counter chart is shown here, indicating that a warning threshold (yellow icon) was breached at 1:00, and that a critical threshold (red icon) was breached at 12:10, 12:30, and 1:10:



A threshold breach must occur continuously for the specified duration. If the threshold dips below the limit values for any reason, a subsequent breach is considered the start of a new duration.

Some cluster objects and performance counters enable you to create a combination threshold policy that requires two performance counters to exceed their maximum limits before an event is generated. For example, you can create a threshold policy using the following criteria:

Cluster object	Performance counter	Warning threshold	Critical threshold	Duration
Volume	Latency	10 milliseconds	20 milliseconds	15 minutes

Threshold policies that use two cluster objects cause an event to be generated only when both conditions are breached. For example, using the threshold policy defined in the table:

If volume latency is averaging	And aggregate disk utilization is	Then
15 milliseconds	50%	No event is reported.
15 milliseconds	75%	A Warning event is reported.
25 milliseconds	75%	A Warning event is reported.
25 milliseconds	90%	A Critical event is reported.

What happens when a performance threshold policy is breached

When a counter value exceeds its defined performance threshold value for the amount of time specified in the duration, the threshold is breached and an event is reported.

The event causes the following actions to be initiated:

- The event is displayed in the Performance Dashboard, the Performance Cluster Summary page, the Events page, and the object-specific Performance Inventory page.
- (optional) An email alert about the event can be sent to one or more email recipients, and an SNMP trap can be sent to a trap receiver.
- (optional) A script can be executed to automatically modify or update storage objects.

The first action is always executed. You configure whether the optional actions are performed in the Configuration/Alerting page page. You can define unique actions depending on whether a Warning or a Critical threshold policy is breached.

After a performance threshold policy breach has occurred on a storage object, no further events are generated for that policy until the counter value goes below the threshold value, at which point the duration resets for that limit. While the threshold continues to be exceeded, the end time of the event is continually updated to reflect that this event is ongoing.

A threshold event captures, or freezes, the information related to severity and policy definition so that unique threshold information displays with the event, even if the threshold policy is modified in the future.

What performance counters can be tracked using thresholds

Some common performance counters, such as IOPS and MBps, can have thresholds set for all storage objects. There are other counters that can have thresholds set for only certain storage objects.

Available performance counters

Storage object	Performance counter	Description
Cluster	IOPS	Average number of input/output operations the cluster processes per second.
MBps	Average number of megabytes of data transferred to and from this cluster per second.	Node
IOPS	Average number of input/output operations the node processes per second.	MBps
Average number of megabytes of data transferred to and from this node per second.	Latency	Average number of milliseconds the node takes to respond to application requests.
Utilization	Average percentage of the node's CPU and RAM that is being used.	Performance Capacity Used
Average percentage of performance capacity that is being consumed by the node.	Performance Capacity Used - Takeover	Average percentage of performance capacity that is being consumed by the node, plus the performance capacity of its partner node.
Aggregate	IOPS	Average number of input/output operations the aggregate processes per second.
MBps	Average number of megabytes of data transferred to and from this aggregate per second.	Latency
Average number of milliseconds the aggregate takes to respond to application requests.	Utilization	Average percentage of the aggregate's disks that are being used.
Performance Capacity Used	Average percentage of performance capacity that is being consumed by the aggregate.	Storage Virtual Machine (SVM)
IOPS	Average number of input/output operations the SVM processes per second.	MBps

Storage object	Performance counter	Description
Average number of megabytes of data transferred to and from this SVM per second.	Latency	Average number of milliseconds the SVM takes to respond to application requests.
Volume	IOPS	Average number of input/output operations the volume processes per second.
MBps	Average number of megabytes of data transferred to and from this volume per second.	Latency
Average number of milliseconds the volume takes to respond to application requests.	Cache miss ratio	Average percentage of read requests from client applications that are returned from the volume instead of being returned from cache.
LUN	IOPS	Average number of input/output operations the LUN processes per second.
MBps	Average number of megabytes of data transferred to and from this LUN per second.	Latency
Average number of milliseconds the LUN takes to respond to application requests.	Namespace	IOPS
Average number of input/output operations the namespace processes per second.	MBps	Average number of megabytes of data transferred to and from this namespace per second.
Latency	Average number of milliseconds the namespace takes to respond to application requests.	Port
Bandwidth utilization	Average percentage of the port's available bandwidth that is being used.	MBps
Average number of megabytes of data transferred to and from this port per second.	Logical Interface (LIF)	MBps



Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

What objects and counters can be used in combination threshold policies

Only some performance counters can be used together in combination policies. When primary and secondary performance counters are specified, both performance counters must exceed their maximum limits before an event is generated.

Primary storage object and counter	Secondary storage object and counter
Volume Latency	Volume IOPS
Volume MBps	Aggregate Utilization
Aggregate Performance Capacity Used	Node Utilization
Node Performance Capacity Used	Node Performance Capacity Used - Takeover
LUN Latency	LUN IOPS
LUN MBps	Aggregate Utilization
Aggregate Performance Capacity Used	Node Utilization
Node Performance Capacity Used	Node Performance Capacity Used - Takeover



When a volume combination policy is applied to a FlexGroup volume, instead of to a FlexVol volume, only the "Volume IOPS" and "Volume MBps" attributes can be selected as the secondary counter. If the threshold policy contains one of the node or aggregate attributes, then the policy will not be applied to the FlexGroup volume, and you will receive an error message describing this case. This is because FlexGroup volumes can exist on more than one node or aggregate.

Creating user-defined performance threshold policies

You create performance threshold policies for storage objects so that notifications are sent when a performance counter exceeds a specific value. The event notification identifies that the cluster is experiencing a performance issue.

Before you begin

You must have the OnCommand Administrator role.

About this task

You create performance threshold policies by entering the threshold values on the Create Threshold Policy page. You can create new policies by defining all the policy values in this page, or you can make a copy of an

existing policy and change a the values in the copy (called *cloning*).

Valid threshold values are 0.001 through 10,000,000 for numbers, 0.001-100 for percentages, and 0.001-200 for Performance Capacity Used percentages.



The current release provides thresholds that send events when a counter value exceeds the threshold setting. You cannot set thresholds that send events when a counter value falls below a threshold setting.

Steps

1. In the left navigation pane, select **Configuration > Performance Thresholds**.

The Configuration/Performance Thresholds page is displayed.

2. Click the appropriate button depending on whether you want to build a new policy or if you want to clone a similar policy and modify the cloned version.

То	Click
Create a new policy	Create
Clone an existing policy	Select an existing policy and click Clone

The Create Threshold Policy page or Clone Threshold Policy page is displayed.

- Define the threshold policy by specifying the performance counter threshold values you want to set for specific storage objects:
 - a. Select the storage object type and specify a name and description for the policy.
 - b. Select the performance counter to be tracked and specify the limit values that define Warning and Critical events.

You must define at least one Warning or one Critical limit. You do not need to define both types of limits.

c. Select a secondary performance counter, if required, and specify the limit values for Warning and Critical events.

Including a secondary counter requires that both counters exceed the limit values before the threshold is breached and an event is reported. Only certain objects and counters can be configured using a combination policy.

d. Select the duration of time for which the limit values must be breached for an event to be sent.

When cloning an existing policy, you must enter a new name for the policy.

4. Click **Save** to save the policy.

You are returned to the Configuration/Performance Thresholds page. A success message at the top of the page confirms that the threshold policy was created and provides a link to the Inventory page for that object type so that you can apply the new policy to storage objects immediately.

After you finish

If you want to apply the new threshold policy to storage objects at this time, you can click the **Go to object_type now** link to go to the Inventory page.

Assigning performance threshold policies to storage objects

You assign a user-defined performance threshold policy to a storage object so that Unified Manager reports an event if the value of the performance counter exceeds the policy setting.

Before you begin

You must have the OnCommand Administrator role.

The performance threshold policy, or policies, that you want to apply to the object must exist.

About this task

You can apply only one performance policy at a time to an object, or to a group of objects.

You can assign a maximum of three threshold policies to each storage object. When assigning policies to multiple objects, if any of the objects already has the maximum number of policies assigned, Unified Manager performs the following actions:

- Applies the policy to all of the selected objects that have not reached their maximum
- Ignores the objects that have reached the maximum number of policies
- Displays a message that the policy was not assigned to all objects

Additionally, if some objects do not support the counter being tracked in the threshold policy, the policy is not applied to that object. For example, if you create a "Performance Capacity Used" threshold policy, and then you attempt to assign it to a node that does not have ONTAP 9.0 or later software installed, the policy is not applied to that node.

Steps

1. From the Performance inventory page of any storage object, select the object or objects to which you want to assign a threshold policy:

To assign thresholds to	Click
A single object	The check box at the left of that object.
Multiple objects	The check box at the left of each object.
All objects on the page	The drop-down box, and choose Select all objects on this page.
All objects of the same type	The drop-down box, and choose Select all objects.

You can use the sorting and filtering functionality to refine the list of objects on the inventory page to make it easier to apply threshold policies to many objects.

2. Make your selection, and then click Assign Performance Threshold Policy.

The Assign Threshold Policy page is displayed, showing a list of threshold policies that exist for that specific type of storage object.

- 3. Click each policy to display the details of the performance threshold settings to verify that you have selected the correct threshold policy.
- 4. After you have selected the appropriate threshold policy, click **Assign Policy**.

A success message at the top of the page confirms that the threshold policy was assigned to the object or objects, and provides a link to the Alerting page so that you can configure alert settings for this object and policy.

After you finish

If you want to have alerts sent over email, or as an SNMP trap, to notify you that a particular performance event has been generated, you must configure the alert settings in the Configuration/Alerting page.

Viewing performance threshold policies

You can view all of the currently defined performance threshold policies from the Configuration/Performance Thresholds page.

About this task

The list of threshold policies is sorted alphabetically by policy name, and it includes policies for all types of storage objects. You can click a column header to sort the policies by that column. If you are looking for a specific policy, use the filter and search mechanisms to refine the list of threshold policies that appear in the inventory list.

You can hover your cursor over the Policy Name and the Condition name to see the configuration details of the policy. Additionally, you can use the provided buttons to create, clone, edit, and delete user-defined threshold policies.

Steps

1. In the left navigation pane, select **Configuration > Performance Thresholds**.

The Configuration/Performance Thresholds page is displayed.

Editing user-defined performance threshold policies

You can edit the threshold settings for existing performance threshold policies. This can be useful if you find that you are receiving too many or too few alerts for certain threshold conditions.

Before you begin

You must have the OnCommand Administrator role.

About this task

You cannot change the policy name or the type of storage object that is being monitored for existing threshold policies.

Steps

1. In the left navigation pane, select Configuration > Performance Thresholds.

The Configuration/Performance Thresholds page displays.

2. Select the threshold policy that you want to change and click Edit.

The Edit Threshold Policy page is displayed.

3. Make your changes to the threshold policy and click **Save**.

You are returned to the Configuration/Performance Thresholds page.

Results

After they are saved, changes are updated immediately on all storage objects that use the policy.

After you finish

Depending on the type of changes that you made to the policy, you may want to review the alert settings configured for the objects that use the policy in the Configuration/Alerting page.

Removing performance threshold policies from storage objects

You can remove a user-defined performance threshold policy from a storage object when you no longer want Unified Manager to monitor the value of the performance counter.

Before you begin

You must have the OnCommand Administrator role.

About this task

You can remove only one policy at a time from a selected object.

You can remove a threshold policy from multiple storage objects by selecting more than one object in the list.

Steps

1. From the **inventory** page of any storage object, select one or more objects that have at least one performance threshold policy applied.

To clear thresholds from	Do this
A single object	Select the check box at the left of that object.
Multiple objects	Select the check box at the left of each object.
All objects on the page	Click and select Select all objects on this page.
All objects of the same type	Click □- and select Select all objects .

2. Click Clear Performance Threshold Policy.

The Clear Threshold Policy page displays, showing a list of threshold policies that are currently assigned to the storage objects.

3. Select the threshold policy you want to remove from the objects and click **Clear Policy**.

When you select a threshold policy, the details of the policy display so that you can confirm that you have selected the appropriate policy.

What happens when a performance threshold policy is changed

If you adjust the counter value or duration of an existing performance threshold policy, the policy change is applied to all storage objects that use the policy. The new setting takes place immediately, and Unified Manager begins to compare performance counter values to the new threshold settings for all newly collected performance data.

If any active events exist for objects that are using the changed threshold policy, the events are marked as obsolete, and the threshold policy begins monitoring the counter as a newly defined threshold policy.

When viewing the counter on which the threshold has been applied in the Counter Charts Detailed View, the critical and warning threshold lines reflect the current threshold settings. The original threshold settings do not appear on this page even if you view historical data when the old threshold setting was in effect.



Because older threshold settings do not appear in the Counter Charts Detailed View, you might see historical events that appear below the current threshold lines.

What happens to performance threshold policies when an object is moved

Because performance threshold policies are assigned to storage objects, if you move an object, all assigned threshold policies remain attached to the object after the move is completed. For example, if you move a volume or LUN to a different aggregate, the threshold policies are still active for the volume or LUN on the new aggregate.

If a secondary counter condition exists for the threshold policy (a combination policy)--for example, if an additional condition is assigned to an aggregate or a node—the secondary counter condition is applied to the new aggregate or node to which the volume or LUN has been moved.

If any new active events exist for objects that are using the changed threshold policy, the events are marked as obsolete, and the threshold policy begins monitoring the counter as a newly defined threshold policy.

A volume move operation causes ONTAP to send an informational change event. A change event icon appears in the Events timeline on the Performance Explorer page and the Performance/Volume Details page to indicate the time when the move operation was completed.



If you move an object to a different cluster, the user-defined threshold policy is removed from the object. If required, you must assign a threshold policy to the object after the move operation is completed. Dynamic and system-defined threshold policies, however, are applied automatically to an object after it has moved to a new cluster.

Threshold policy functionality during HA takeover and giveback

When a takeover or giveback operation occurs in a high-availability (HA) configuration, objects that are moved from one node to the other node retain their threshold policies in the same manner as in the manual move operations. Because Unified Manager checks for cluster configuration changes every 15 minutes, the impact of the switchover to the new node is not identified until the next poll of the cluster configuration.



If both a takeover and giveback operation occur within the 15-minute configuration change collection period, you might not see the performance statistics move from one node to the other node.

Threshold policy functionality during aggregate relocation

If you move an aggregate from one node to another node using the aggregate relocation start command, both single and combination threshold policies are retained on all objects, and the node portion of the threshold policy is applied to the new node.

Threshold policy functionality during MetroCluster switchover

Objects that move from one cluster to another cluster in a MetroCluster configuration do not retain their user-defined threshold policy settings. If required, you can apply threshold policies on the volumes and LUNs that have moved to the partner cluster. After an object has moved back to its original cluster, the user-defined threshold policy is reapplied automatically.

Volume behavior during switchover and switchback

Descriptions of the performance threshold policy pages

You use the Configuration/Performance Thresholds page to create, edit, clone, delete, and view performance threshold policies.

The topics below display when you click **Help** on the appropriate page.

Configuration/Performance Thresholds page

You can use the Configuration/Performance Thresholds page to view all the currently defined performance threshold policies. This page also provides the functionality to create, clone, edit, and delete threshold policies.

The list of performance threshold policies is sorted alphabetically by policy name. You can click a column

header to sort the policies by that column. If you are looking for a specific policy, you can use the filter and search mechanisms to refine the list of threshold policies that appear in the inventory list.

Filter and Search bar

The **Filtering** button enables you to refine the list of threshold policies by displaying only the policies that match certain criteria.

The **Search** button enables you to search for certain policies by entering full or partial policy names to refine the list of threshold policies that appear in the inventory list.

Command buttons

Create

Creates a new performance threshold policy.

Clone

Creates a new performance threshold policy based on a copy of the policy that you have selected.

• Edit

Modifies the performance threshold policy that you have selected. All storage objects that are using the policy are updated to use the revised policy.

Delete

Deletes the performance threshold policy that you have selected. The policy is removed from all storage objects that are using the policy. You can click the item in the Associated Objects column to view the objects that are currently using this policy.

Threshold Policies list

Policy Name

Displays the name of the threshold policy. You can position your cursor over the policy name to view the details of the policy.

Description

Displays a brief description of the threshold policy.

First Condition

Displays the primary condition for the threshold policy, including the defined performance counter and the warning trigger values and critical trigger values. You can position your cursor over the condition name to view the details of the condition.

Second Condition

Displays the secondary threshold policy condition, if defined. You can position your cursor over the condition name to view the details of the condition. If a second condition is not defined, this column is blank.



When a second condition is defined, an event is generated only when both conditions are breached.

Associated Objects

Displays the type of storage object to which the threshold policy can be applied, and the number of objects that are using the policy. This field is blank until you assign the policy to at least one object.

You can click the column heading to sort the policies by object type: volume, LUN, aggregate, and so on. You can click the policy name to display the inventory page populated with the objects that are currently using the threshold policy.

Create or Clone Performance Threshold Policy page

You can use the Create Threshold Policy page or the Clone Threshold Policy page to create a new performance threshold policy.

You can complete the fields on this page and click **Save** to add a performance threshold policy.

For Object Type

Select the type of storage object for which you want to create a threshold policy.

Policy Name

Enter the name of the threshold policy. The name appears on other Unified Manager pages and should provide a brief description of the policy.

Description

(optional) Enter a detailed description of the threshold policy.

Threshold Values

Define the primary, and optionally the secondary, threshold counter condition. Including a secondary counter requires that both counters exceed the limit values before the threshold is considered breached.

Select a counter

Select the counter on which you want to set a performance threshold.

Warning

Enter the limit value for the counter that is considered a warning.

Critical

Enter the limit value for the counter that is considered critical.

Valid threshold values are 0.001 through 10,000,000 for numbers, 0.001-100 for percentages, and 0.001-200 for Performance Capacity Used percentages.

Duration

Select the number of minutes that the counter value must be greater than the warning or critical limit value. Because Unified Manager collects new performance counter values every five minutes, the menu provides values in multiples of five based on the refresh interval.

Edit Performance Threshold Policy page

You can use the Edit Threshold Policy page to modify an existing performance threshold policy.

You can modify the fields on this page and click **Save** to change a performance threshold policy. All cluster objects that are currently using the threshold policy are automatically updated to use the new policy definition.

For Object Type

Object type cannot be changed.

Policy Name

Change the name of the threshold policy.

Description

Change the detailed description of the threshold policy.

Threshold Values

Change the primary, and optionally the secondary, threshold counter condition.

Select a counter

Change the counter on which you want to set a performance threshold.

Warning

Enter the limit value for the counter that is considered a warning.

Critical

Enter the limit value for the counter that is considered critical.

Duration

Change the number of minutes that the counter value must be greater than the warning or critical limit value.

Assign Performance Threshold Policy page

You can use the Assign Threshold Policy page to assign a performance threshold policy to one or more storage objects.

The list of policies is populated with only those policies that are valid for the storage object type.

You select the that policy you want to apply to the object or objects, and then you click Apply Policy.

There are a few cases where an error message may be returned when you attempt to apply a policy:

• When applying a policy that uses the Performance Capacity Used counter to a node or aggregate that is not installed with ONTAP 9.0, or later, software.

Versions of ONTAP software prior to 9.0 do not support the performance capacity counters.

• When applying a combination policy to a FlexGroup volume, where the second counter includes either a node or aggregate object.

Because FlexGroup volumes can be spread across multiple nodes and aggregates, this operation is not allowed.

Clear Performance Threshold Policy page

You can use the Clear Threshold Policy page to remove, or *clear*, a performance threshold policy from one or more storage objects.

The list of policies is populated with only those policies that are being used in the selected object or objects.

You select the policy that you want to remove from the storage object or objects, and then you click **Clear Policy**.

Analyzing performance events

You can analyze performance events to identify when they were detected, whether they are active (new or acknowledged) or obsolete, the workloads and cluster components involved, and the options for resolving the events on your own.

Displaying information about performance events

You can use the Events inventory page to view a list of all the new and obsolete performance events on the clusters being monitored by Unified Manager. By viewing this information you can determine the most critical events and then drill down to detailed information to determine the cause of the event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

About this task

The list of events is sorted by detected time, with the most recent events listed first. You can click a column header to sort the events based on that column. For example, you can sort by the Status column to view events by severity. If you are looking for a specific event, or for a specific type of event, you can use the filter and search mechanisms to refine the list of events that appear in the list.

Events from all sources are displayed on this page:

User-defined performance threshold policy

- System-defined performance threshold policy
- · Dynamic performance threshold

The Event Type column lists the source of the event. You can select an event to view details about the event in the Event details page.

Steps

- 1. In the left navigation pane, click **Events**.
- 2. Locate an event that you want to analyze and click the event name.

The details page for the event displays.



You can also display the details page for an event by clicking the event name link from the Performance Explorer page and from an alert email.

Analyzing events from user-defined performance thresholds

Events generated from user-defined thresholds indicate that a performance counter for a certain storage object, for example, an aggregate or volume, has crossed the threshold you defined in the policy. This indicates that the cluster object is experiencing a performance issue.

You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.

Responding to user-defined performance threshold events

You can use Unified Manager to investigate performance events caused by a performance counter crossing a user-defined warning or critical threshold. You can also use Unified Manager to check the health of the cluster component to see whether recent health events detected on the component contributed to the performance event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

Steps

- 1. Display the **Event** details page to view information about the event.
- 2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message "Latency value of 456 ms/op has triggered a WARNING event based on threshold setting of 400 ms/op" indicates that a latency warning event occurred for the object.

3. Hover your cursor over the policy name to display details about the threshold policy that triggered the event.

This includes the policy name, the performance counter being evaluated, the counter value that must be

breached to be considered a critical or warning event, and the duration by which the counter must exceed the value.

- 4. Make a note of the **Event Trigger Time** so you can investigate whether other events might have occurred at the same time that could have contributed to this event.
- 5. Follow one of the options below to further investigate the event, to determine whether you need to perform any actions to resolve the performance problem:

Option	Possible investigation actions			
Click the Source object name to display the Explorer page for that object.	This page enables you to view the object details and compare this object with other similar storage objects to see whether other storage objects have a performance issue around the same time. For example, to see whether other volumes on the same aggregate are also having a performance issue.			
Click the cluster name to display the Cluster Summary page.	This page enables you to view the details for the cluster on which this object resides to see whether other performance issues have occurred around the same time.			

Analyzing events from system-defined performance thresholds

Events generated from system-defined performance thresholds indicate that a performance counter, or set of performance counters, for a certain storage object has crossed the threshold from a system-defined policy. This indicates that the storage object, for example, an aggregate or node, is experiencing a performance issue.

You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.



System-defined threshold policies are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

Responding to system-defined performance threshold events

You can use Unified Manager to investigate performance events caused by a performance counter crossing a system-defined warning threshold. You can also use Unified Manager to check the health of the cluster component to see whether recent events detected on the component contributed to the performance event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

Steps

- 1. Display the **Event** details page to view information about the event.
- 2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message "Node utilization value of 90 % has triggered a WARNING event based on threshold setting of 85 %" indicates that a node utilization warning event occurred for the cluster object.

- 3. Make a note of the **Event Trigger Time** so you can investigate whether other events might have occurred at the same time that could have contributed to this event.
- 4. Under **System Diagnosis**, review the brief description of the type of analysis the system-defined policy is performing on the cluster object.

For some events a green or red icon is displayed next to the diagnosis to indicate whether an issue was found in that particular diagnosis. For other types of system-defined events counter charts display the performance for the object.

5. Under **Suggested Actions**, click the **Help me do this** link to view the suggested actions you can perform to try and resolve the performance event on your own.

Responding to QoS policy group performance events

Unified Manager generates QoS policy warning events when workload throughput (IOPS, IOPS/TB, or MBps) has exceeded the defined ONTAP QoS policy setting and workload latency is becoming affected. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

About this task

Unified Manager generates warning events for QoS policy breaches when workload throughput has exceeded the defined QoS policy setting during each performance collection period for the previous hour. Workload throughput may exceed the QoS threshold for only a short period of time during each collection period, but Unified Manager displays only the "average" throughput during the collection period on the chart. For this reason you may receive QoS events while the throughput for a workload might not have crossed the policy threshold shown in the chart.

You can use System Manager or the ONTAP commands to manage policy groups, including the following tasks:

- · Creating a new policy group for the workload
- Adding or removing workloads in a policy group
- · Moving a workload between policy groups
- · Changing the throughput limit of a policy group
- · Moving a workload to a different aggregate or node

Steps

- 1. Display the **Event** details page to view information about the event.
- 2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message "IOPS value of 1,352 IOPS on vol1_NFS1 has triggered a WARNING event to identify potential performance problems for the workload" indicates that a QoS Max IOPS event occurred on volume vol1_NFS1.

3. Review the **Event Information** section to see more details about when the event occurred and how long the event has been active.

Additionally, for volumes or LUNs that are sharing the throughput of a QoS policy you can see the names of the top three workloads that are consuming the most IOPS or MBps.

4. Under the System Diagnosis section, review the two charts: one for total average IOPS or MBps (depending on the event), and one for latency. When arranged this way you can see which cluster components are most affecting latency when the workload approached the QoS max limit.

For a shared QoS policy event, the top three workloads are shown in the throughput chart. If more than three workloads are sharing the QoS policy, then additional workloads are added together in an "Other workloads" category. Additionally, the Latency chart shows the average latency on all workloads that are part of the QoS policy.

Note that for adaptive QoS policy events that the IOPS and MBps charts show IOPS or MBps values that ONTAP has converted from the assigned IOPS/TB threshold policy based on the size of the volume.

5. Under the **Suggested Actions** section, review the suggestions and determine which actions you should perform to avoid an increase in latency for the workload.

If required, click the **Help** button to view more details about the suggested actions you can perform to try and resolve the performance event.

Understanding events from adaptive QoS policies that have a defined block size

Adaptive QoS policy groups automatically scale a throughput ceiling or floor based on the volume size, maintaining the ratio of IOPS to TBs as the size of the volume changes. Starting with ONTAP 9.5 you can specify the block size in the QoS policy to effectively apply a MBps threshold at the same time.

Assigning an IOPS threshold in an adaptive QoS policy places a limit only on the number of operations that occur in each workload. Depending on the block size that is set on the client that generates the workloads, some IOPS include much more data and therefore place a much larger burden on the nodes that process the operations.

The MBps value for a workload is generated using the following formula:

```
MBps = (IOPS * Block Size) / 1000
```

If a workload is averaging 3,000 IOPS and the block size on the client is set to 32 KB, then the effective MBps for this workload is 96. If this same workload is averaging 3,000 IOPS and the block size on the client is set to 48 KB, then the effective MBps for this workload is 144. You can see that the node is processing 50% more

data when the block size is larger.

Let's look at the following adaptive QoS policy that has a defined block size and how events are triggered based on the block size that is set on the client.

Create a policy and set the peak throughput to 2,500 IOPS/TB with a block size of 32KB. This effectively sets the MBps threshold to 80 MBps ((2500 IOPS * 32KB) / 1000) for a volume with 1 TB used capacity. Note that Unified Manager generates a Warning event when the throughput value is 10% less than the defined threshold. Events are generated under the following situations:

Used Capacity	Event is generated when throughput exceeds this number of
IOPS	MBps
1 TB	2,250 IOPS
72 MBps	2 TB
4,500 IOPS	144 MBps
5 TB	11,250 IOPS

If the volume is using 2TB of the available space, and the IOPS is 4,000, and the QoS block size is set to 32KB on the client, then the MBps throughput is 128 MBps ((4,000 IOPS * 32 KB) / 1000). No event is generated in this scenario because both 4,000 IOPS and 128 MBps are below the threshold for a volume that is using 2 TB of space.

If the volume is using 2TB of the available space, and the IOPS is 4,000, and the QoS block size is set to 64KB on the client, then the MBps throughput is 256 MBps ((4,000 IOPS * 64 KB) / 1000). In this case the 4,000 IOPS does not generate an event, but the MBps value of 256 MBps is above the threshold of 144 MBps and an event is generated.

For this reason, when an event is triggered based on a MBps breach for an adaptive QoS policy that includes the block size, a MBps chart is displayed in the System Diagnosis section of the Event details page. If the event is triggered based on an IOPS breach for the adaptive QoS policy, an IOPS chart is displayed in the System Diagnosis section. If a breach occurs for both IOPS and MBps you will receive two events.

For more information on adjusting QoS settings, see the ONTAP 9 Performance Monitoring Power Guide.

ONTAP 9 Performance Monitoring Power Guide

Responding to node resources overutilized performance events

Unified Manager generates node resources overutilized warning events when a single node is operating above the bounds of its operational efficiency, and therefore potentially affecting workload latencies. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

Before you begin

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

• There must be new or obsolete performance events.

About this task

Unified Manager generates warning events for node resources overutilized policy breaches by looking for nodes that are using more than 100% of their performance capacity for more than 30 minutes.

You can use System Manager or the ONTAP commands to correct this type of performance issue, including the following tasks:

- Creating and applying a QoS policy to any volumes or LUNs that are overusing system resources
- Reducing the QoS maximum throughput limit of a policy group to which workloads have been applied
- · Moving a workload to a different aggregate or node
- Increasing capacity by adding disks to the node, or by upgrading to a node with a faster CPU and more RAM

Steps

- 1. Display the **Event** details page to view information about the event.
- 2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message "Perf. Capacity Used value of 139% on simplicity-02 has triggered a WARNING event to identify potential performance problems in the data processing unit." indicates that performance capacity on node simplicity-02 is overused and affecting node performance.

3. Under the System Diagnosis section, review the three charts: one for performance capacity used on the node, one for average storage IOPS being used by the top workloads, and one for latency on the top workloads. When arranged in this way you can see which workloads are the cause of the latency on the node.

You can view which workloads have QoS policies applied, and which do not, by moving your cursor over the IOPS chart.

4. Under the **Suggested Actions** section, review the suggestions and determine which actions you should perform to avoid an increase in latency for the workload.

If required, click the **Help** button to view more details about the suggested actions you can perform to try and resolve the performance event.

Analyzing events from dynamic performance thresholds

Events generated from dynamic thresholds indicate that the actual response time (latency) for a workload is too high, or too low, compared to the expected response time range. You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.



Dynamic performance thresholds are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

Identifying victim workloads involved in a dynamic performance event

In Unified Manager, you can identify which volume workloads have the highest deviation in response time (latency) caused by a storage component in contention. Identifying these workloads helps you understand why the client applications accessing them have been performing slower than usual.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete dynamic performance events.

About this task

The Event details page displays a list of the user-defined and system-defined workloads, ranked by the highest deviation in activity or usage on the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. In the Workload Latency and Workload Activity charts, select Victim Workloads.
- 3. Hover your cursor over the charts to view the top user-defined workloads that are affecting the component, and the name of the victim workload.

Identifying bully workloads involved in a dynamic performance event

In Unified Manager, you can identify which workloads have the highest deviation in usage for a cluster component in contention. Identifying these workloads helps you understand why certain volumes on the cluster have slow response times (latency).

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete dynamic performance events.

About this task

The Event details page displays a list of the user-defined and system-defined workloads ranked by the highest usage of the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. In the Workload Latency and Workload Activity charts, select Bully Workloads .
- 3. Hover your cursor over the charts to view the top user-defined bully workloads that are affecting the component.

Identifying shark workloads involved in a dynamic performance event

In Unified Manager, you can identify which workloads have the highest deviation in usage for a storage component in contention. Identifying these workloads helps you determine if these workloads should be moved to a less-utilized cluster.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There are new, acknowledged, or obsolete performance dynamic event.

About this task

The Event details page displays a list of the user-defined and system-defined workloads ranked by the highest usage of the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. In the Workload Latency and Workload Activity charts, select **Shark Workloads**.
- 3. Hover your cursor over the charts to view the top user-defined workloads that are affecting the component, and the name of the shark workload.

Performance event analysis for a MetroCluster configuration

You can use Unified Manager to analyze a performance event for a MetroCluster configuration. You can identify the workloads involved in the event and review the suggested actions for resolving it.

MetroCluster performance events might be due to *bully* workloads that are over-utilizing the interswitch links (ISLs) between the clusters, or due to link health issues. Unified Manager monitors each cluster in a MetroCluster configuration independently, without consideration of performance events on a partner cluster.

Performance events from both clusters in the MetroCluster configuration are also displayed on the Unified ManagerDashboards/Overview page. You can also view the Health pages of Unified Manager to check the health of each cluster and to view their relationship.

Analyzing a dynamic performance event on a cluster in a MetroCluster configuration

You can use Unified Manager to analyze the cluster in a MetroCluster configuration on which a performance event was detected. You can identify the cluster name, event detection time, and the *bully* and *victim* workloads involved.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events for a MetroCluster configuration.
- Both clusters in the MetroCluster configuration must be monitored by the same instance of Unified Manager.

Steps

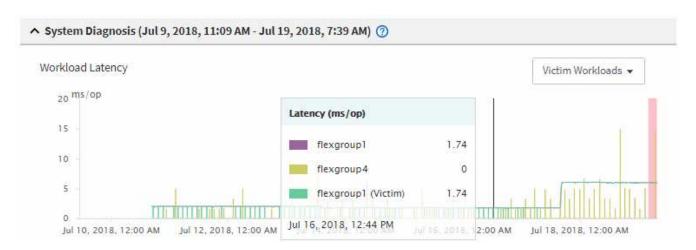
- 1. Display the **Event details** page to view information about the event.
- 2. Review the event description to see the names of the workloads involved and the number of workloads involved.

In this example, the MetroCluster Resources icon is red, indicating that the MetroCluster resources are in contention. You position your cursor over the icon to display a description of the icon. At the top of the page in the event ID, the cluster name identifies the name of the cluster on which the event was detected.



- 3. Make a note of the cluster name and the event detection time, which you can use to analyze performance events on the partner cluster.
- 4. In the charts, review the *victim* workloads to confirm that their response times are higher than the performance threshold.

In this example, the victim workload is displayed in the hover text. The Latency charts display, at a high-level, a consistent latency pattern for the victim workloads involved. Even though the abnormal latency of the victim workloads triggered the event, a consistent latency pattern might indicate that the workloads are performing within their expected range, but that a spike in I/O increased the latency and triggered the event.



If you recently installed an application on a client that accesses these volume workloads and that application sends a high amount of I/O to them, you might be anticipating their latencies to increase. If the latency for the workloads returns within the expected range, the event state changes to obsolete, and remains in this state for more than 30 minutes, you can probably ignore the event. If the event is ongoing, and remains in the new state, you can investigate it further to determine whether other issues caused the event.

In the Workload Throughput chart, select Bully Workloads to display the bully workloads.

The presence of bully workloads indicates that the event might have been caused by one or more workloads on the local cluster overutilizing the MetroCluster resources. The bully workloads have a high

deviation in write throughput (MBps).

This chart displays, at a high-level, the write throughput (MBps) pattern for the workloads. You can review the write MBps pattern to identify abnormal throughput, which might indicate that a workload is over-utilizing the MetroCluster resources.

If no bully workloads are involved in the event, the event might have been caused by a health issue with the link between the clusters or a performance issue on the partner cluster. You can use Unified Manager to check the health of both clusters in a MetroCluster configuration. You can also use Unified Manager to check for and analyze performance events on the partner cluster.

Analyzing a dynamic performance event for a remote cluster on a MetroCluster configuration

You can use Unified Manager to analyze dynamic performance events on a remote cluster in a MetroCluster configuration. The analysis helps you determine whether an event on the remote cluster caused an event on its partner cluster.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have analyzed a performance event on a local cluster in a MetroCluster configuration and obtained the event detection time.
- You must have checked the health of the local cluster and its partner cluster involved in the performance event and obtained the name of the partner cluster.

Steps

- 1. Log in to the Unified Manager instance that is monitoring the partner cluster.
- 2. In the left navigation pane, click **Events** to display the event list.
- 3. From the **Time Range** selector, select **Last Hour**, and then click **Apply Range**.
- 4. In the **Filtering** selector, select **Cluster** from the left drop-down menu, type the name of the partner cluster in the text field, and then click **Apply Filter**.
 - If there are no events for the selected cluster over the last hour, this indicates that the cluster has not experienced any performance issues during the time that the event was detected on its partner.
- 5. If the selected cluster has events detected over the last hour, compare the event detection time to the event detection time for the event on the local cluster.
 - If these events involve bully workloads causing contention on the data processing component, one or more of these bullies might have caused the event on the local cluster. You can click the event to analyze it and review the suggested actions for resolving it on the Event details page.

If these events do not involve bully workloads, they did not cause the performance event on the local cluster.

Responding to a dynamic performance event caused by QoS policy group throttling

You can use Unified Manager to investigate a performance event caused by a Quality of Service (QoS) policy group throttling workload throughput (MBps). The throttling

increased the response times (latency) of volume workloads in the policy group. You can use the event information to determine whether new limits on the policy groups are needed to stop the throttling.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. Read the **Description**, which displays the name of the workloads impacted by the throttling.



The description can display the same workload for the victim and bully, because the throttling makes the workload a victim of itself.

3. Record the name of the volume, using an application such as a text editor.

You can search on the volume name to locate it later.

- 4. In the Workload Latency and Workload Activity charts, select **Bully Workloads**.
- 5. Hover your cursor over the charts to view the top user-defined workloads that are affecting the policy group.

The 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.

- 6. Navigate to the **Performance/Volume Details** page for the top workload.
- 7. Select Break down data by.
- 8. Select the check box next to **Latency** to select all latency breakdown charts.
- 9. Under IOPS, select Reads/writes/other.
- 10. Click Submit.

The breakdown charts are displayed under the Latency chart and the IOPS chart.

11. Compare the **Policy Group Impact** chart to the **Latency** chart to see what percentage of throttling impacted the latency at the time of the event.

The policy group has a maximum throughput of 1,000 operations per second (op/sec), which the workloads in it cannot collectively exceed. At the time of the event, the workloads in the policy group had a combined throughput of over 1,200 op/sec, which caused the policy group to throttle its activity back to 1,000 op/sec. The Policy Group Impact chart shows that the throttling caused 10% of the total latency, confirming that the throttling caused the event to occur.

12. Review the Cluster Components chart, which shows the total latency by cluster component.

The latency is highest at the policy group, further confirming that the throttling caused the event.

13. Compare the **Reads/writes latency** 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 there is a high amount of throughput or number of operations that increased the latency. You can use these values when deciding to put a policy group limit on the throughput or operations.

- 14. Use OnCommand System Manager to increase the current limit on the policy group to 1,300 op/sec.
- 15. After a day, return to Unified Manager and search for the name of the workload that you recorded in Step 3.

The Performance/Volume Details page is displayed.

- 16. Select Break down data by > IOPS.
- 17. Click Submit.

The Reads/writes/other chart is displayed.

- 18. At the bottom of the page, point your cursor to the change event icon () for the policy group limit change.
- 19. Compare the **Reads/writes/other** chart to the **Latency** chart.

The read and write requests are the same, but the throttling has stopped and the latency has decreased.

Responding to a dynamic performance event caused by a disk failure

You can use Unified Manager to investigate a performance event caused by workloads overutilizing an aggregate. You can also use Unified Manager to check the health of the aggregate to see if recent health events detected on the aggregate contributed to the performance event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. Read the **Description**, which describes the workloads involved in the event and the cluster component in contention.

There are multiple victim volumes whose latency was impacted by the cluster component in contention. The aggregate, which is in the middle of a RAID reconstruct to replace the failed disk with a spare disk, is the cluster component in contention. Under Component in Contention, the Aggregate icon is highlighted red and the name of the aggregate is displayed in parentheses.

- 3. In the Workload Utilization chart, select **Bully Workloads**.
- 4. Hover your cursor over the chart to view the top bully workloads that are affecting the component.

The top workloads with the highest peak utilization since the event was detected are displayed at the top of the chart. One of the top workloads is the system-defined workload Disk Health, which indicates a RAID reconstruct. A reconstruct is the internal process involved with rebuilding the aggregate with the spare disk. The Disk Health workload, along with other workloads on the aggregate, likely caused the contention on

the aggregate and the associated event.

- After confirming that the activity from the Disk Health workload caused the event, wait for approximately 30 minutes for the reconstruction to finish and for Unified Manager to analyze the event and detect whether the aggregate is still in contention.
- 6. In Unified Manager, search for the event ID you recorded in Step 2.

The event for the disk failure is displayed on the Event details page. After the RAID reconstruction is complete, check that the State is obsolete, indicating that the event is resolved.

- 7. In the Workload Utilization chart, select **Bully Workloads** to view the workloads on the aggregate by peak utilization.
- 8. Navigate to the **Performance/Volume Details** page for the top workload.
- 9. Click 1d to display the last 24 hours (1 day) of data for the selected volume.

In the Latency chart, a red dot () indicates when the disk failure event occurred.

- 10. Select Break down data by.
- 11. Under Components, select Disk Utilization.
- 12. Click Submit.

The Disk Utilization chart displays a graph of all read and write requests from the selected workload to the disks of the target aggregate.

13. Compare the data in the **Disk Utilization** chart to the data at the time of the event in the **Latency** chart.

At the time of the event, the Disk Utilization shows a high amount of read and write activity, caused by the RAID reconstruction processes, which increased the latency of the selected volume. A few hours after the event occurred, both the reads and writes and the latency have decreased, confirming that the aggregate is no longer in contention.

Responding to a dynamic performance event caused by HA takeover

You can use Unified Manager to investigate a performance event caused by high data processing on a cluster node that is in a high-availability (HA) pair. You can also use Unified Manager to check the health of the nodes to see whether any recent health events detected on the nodes contributed to the performance event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. Read the **Description**, which describes the workloads involved in the event and the cluster component in contention.

There is one victim volume whose latency was impacted by the cluster component in contention. The data processing node, which took over all workloads from its partner node, is the cluster component in

contention. Under Component in Contention, the Data Processing icon is highlighted red and the name of the node that was handling data processing at the time of the event is displayed in parentheses.

3. In the **Description**, click the name of the victim volume.

The Performance/Volume Details page is displayed. At the bottom of the page, in the Events time line, a change event icon () indicates the time that Unified Manager detected the start of the HA takeover.

4. Point your cursor to the change event icon for the HA takeover.

Details about the HA takeover are displayed in the Events List table. In the Latency chart, an event indicates that the selected volume crossed the performance threshold due to high latency around the same time as the HA takeover.

- Select Break down data by.
- 6. Under Latency, select Cluster Components.
- 7. Click Submit.

The Cluster Components chart is displayed. The chart breaks down the total latency by cluster component.

- 8. At the bottom of the page, point your mouse cursor to the change event icon for the start of the HA takeover.
- In the Cluster Components chart, compare the latency for data processing to the total latency in the Latency chart.

At the time of the HA takeover, there was a spike in data processing from the increased workload demand on the data processing node. The increased CPU utilization drove up the latency and triggered the event.

- 10. After fixing the failed node, use OnCommand System Manager to perform an HA giveback, which moves the workloads from the partner node to the fixed node.
- 11. After the HA giveback is complete, in Unified Manager, search for the event ID you recorded in Step 2.

The event triggered by the HA takeover is displayed on the Event details page. The event now has a state of obsolete, which indicates that the event is resolved.

12. In the **Description**, click the name of the victim volume.

The Performance/Volume Details page is displayed. At the bottom of the page, in the Events time line, a change event icon indicates the time that Unified Manager detected the completion of the HA giveback.

- 13. Select Break down data by.
- 14. Under Latency, select Cluster Components.

The Cluster Components chart is displayed.

15. At the bottom of the page, point your cursor to the change event icon for the HA giveback.

The change event is highlighted in the Events List table and indicates that the HA giveback was completed successfully.

16. In the **Cluster Components** chart, compare the latency for data processing to the total latency in the **Latency** chart.

The latency at the data processing component has decreased, which has decreased the total latency. The node that the selected volume is now using for data processing has resolved the event.

Resolving performance events

You can use the suggested actions to try and resolve performance events on your own. The first three suggestions are always displayed, and the actions under the fourth suggestion are specific to the type of event displayed.

The **Help me do this** links provide additional information for each suggested action, including instructions for performing a specific action. Some of the actions may involve using Unified Manager, OnCommand System Manager, OnCommand Workflow Automation, ONTAP CLI commands, or a combination of these tools.

Confirming that the latency is within the expected range

When a cluster component is in contention, volume workloads that use it might have decreased response time (latency). You can review the latency of each victim workload on the component in contention to confirm that its actual latency is within its expected range. You can also click a volume name to view the historical data for the volume.

If the performance event is in the obsolete state, the latency of each victim involved in the event might have returned back within its expected range.

Review the impact of configuration changes on workload performance

Configuration changes on the cluster, such as a failed disk, HA failover, or a moved volume, could negatively impact volume performance and cause increased latency.

In Unified Manager, you can review the Performance/Volume Details page to see when a recent configuration change occurred and compare it to the operations and latency (response time) to see whether there was a change in activity for the selected volume workload.

The performance pages of Unified Manager can only detect a limited number of change events. The health pages provide alerts for other events caused by configuration changes. You can search for the volume in Unified Manager to see the event history.

Options for improving workload performance from the client-side

You can check your client workloads, such as applications or databases, that are sending I/O to volumes involved in a performance event to determine if a client-side change might correct the event.

When the clients that are connected to volumes on a cluster increase their I/O requests, the cluster must work harder to meet the demand. If you know which clients have a high number of I/O requests to a particular volume on the cluster, you can improve cluster performance by adjusting the number of clients accessing the volume or decreasing the amount of I/O to the volume. You can also apply or increase a limit on the QoS policy group of which the volume is a member.

You can investigate clients and their applications to determine whether the clients are sending more I/O than usual, which might be causing contention on a cluster component. On the Event details page, the System

Diagnosis section displays the top volume workloads using the component in contention. If you know which client is accessing a particular volume, you can go to the client to determine whether the client hardware or an application is not operating as expected or is doing more work than usual.

In a MetroCluster configuration, write requests to a volume on a local cluster are mirrored to a volume on the remote cluster. Keeping the source volume on the local cluster in sync with the destination volume on the remote cluster can also increase the demand of both clusters in the MetroCluster configuration. By reducing write requests to these mirrored volumes, the clusters can perform fewer sync operations, which reduces the performance impact on other workloads.

Check for client or network issues

When the clients that are connected to volumes on a cluster increase their I/O requests, the cluster must work harder to meet the demand. The increased demand on the cluster can put a component in contention, increase the latency of workloads that use it, and trigger an event in Unified Manager.

On the Event details page, the System Diagnosis section displays the top volume workloads using the component in contention. If you know which client is accessing a particular volume, you can go to the client to determine whether the client hardware or an application is not operating as expected or is doing more work than usual. You might need to contact your client administrator or application vendor for assistance.

You can check your network infrastructure to determine whether there are hardware issues, bottlenecks, or competing workloads that might have caused I/O requests between the cluster and connected clients to perform slower than expected. You might need to contact your network administrator for assistance.

Verify whether other volumes in the QoS policy group have unusually high activity

You can review the workloads in the Quality of Service (QoS) policy group with the highest change in activity to determine whether more than one workload caused the event. You can also see whether other workloads are still exceeding the set throughput limit or whether they are back within their expected range of activity.

On the Event details page, in the System Diagnosis section, you can sort the workloads by peak deviation in activity to display the workloads with the highest change in activity at the top of the table. These workloads might be the "bullies" whose activity exceeded the set limit and might have caused the event.

You can navigate to the Performance/Volume Details page for each volume workload in the chart to review its IOPS activity. If the workload has periods of very high operations activity, it might have contributed to the event. You can change the policy group settings for the workload or move the workload to a different policy group.

You can use OnCommand System Manager or the ONTAP CLI commands to manage policy groups, as follows:

- · Create a policy group.
- Add or remove workloads in a policy group.
- · Move a workload between policy groups.
- · Change the throughput limit of a policy group.

Move logical interfaces (LIFs)

Moving logical interfaces (LIFs) to a less busy port can help improve load balancing, assist with maintenance operations and performance tuning, and reduce indirect access.

Indirect access can reduce system efficiency. It occurs when a volume workload is using different nodes for network processing and data processing. To reduce indirect access, you can rearrange LIFs, which involves moving LIFs to use the same node for network processing and data processing. You can configure load balancing to have ONTAP automatically move busy LIFs to a different port or you can move a LIF manually.

Benefits

- · Improve load balancing.
- · Reduce indirect access.

Considerations



When moving a LIF connected to CIFS shares, clients accessing the CIFS shares are disconnected. Any read or write requests to the CIFS shares are disrupted.

You use the ONTAP commands to configure load balancing. For more information, see the ONTAP networking documentation.

You use OnCommand System Manager and the ONTAP CLI commands to move LIFs manually.

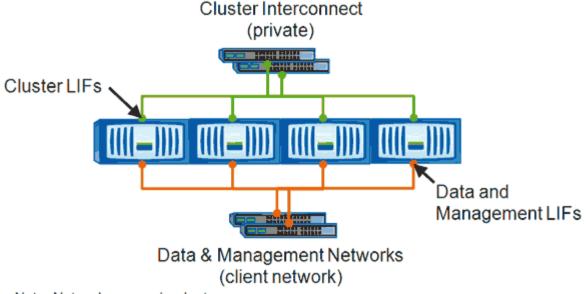
Moving LIFs manually

storage virtual machines (SVMs)contain data volumes and one or more logical interfaces (LIFs) through which the SVM provides data to the clients. You can move data LIFs from one physical port to another within the same SVM. You might want to do this to improve load balancing or assist with maintenance operations and performance tuning.

About this task

The following types of LIFs exist:

- Data LIFs: Associated with a SVM and used for communicating with clients.
- Cluster Management LIFs: Used for managing nodes, SVMs, and the cluster itself.
- · Cluster LIFs: Used for intracluster traffic.
- Intercluster LIFs: Used for communication between clusters.
- Intracluster LIFs: Used for communication between HA pairs.
- SVM Management LIFs: Data LIFs associated with a SVM and used for managing that SVM.



Note: Networks are redundant

This workflow describes how to move data LIFs. This applies to NAS (NFS and CIFS) LIFs, but not to SAN (FC and iSCSI) LIFs.



When moving a LIF connected to CIFS shares, clients accessing the CIFS shares will be disconnected. Any read or write requests to the CIFS shares will be disrupted.



For information about how to move other types of LIFs, including details about moving LIFS connected CIFS shares, see the ONTAP networking documentation.

You can perform the following basic actions related to data LIFs:

- · Display all the data LIFs.
- · Identify the busiest LIFs.
- Identify the best node to accept a busy LIF.
- · Modify the home port or node for a LIF to change its preferred location in the cluster.

You should move a LIF rather than migrate a LIF for a more lasting change. To return to the original home port, you should revert the LIF.

- Migrate a data LIF to another port for a temporary change that might be used if the home port or node has a problem or is undergoing scheduled maintenance.
- · Revert a data LIF to its home port.

What LIFs are

A LIF (logical interface) is an IP address or WWPN with associated characteristics, such as a role, a home port, a home node, a list of ports to fail over to, and a firewall policy. You can configure LIFs on ports over which the cluster sends and receives communications over the network.

LIFs can be hosted on the following ports:

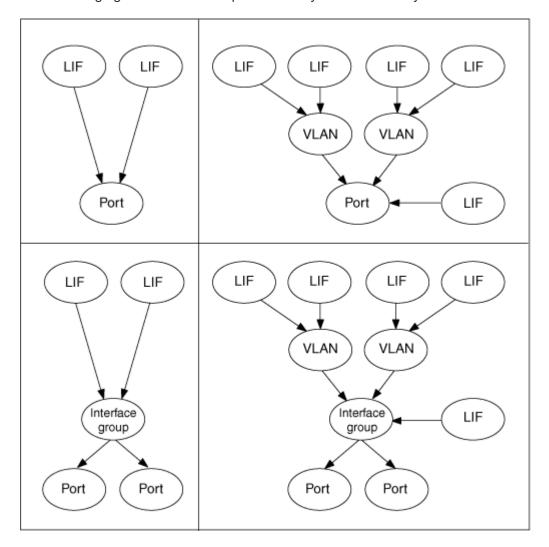
- · Physical ports that are not part of interface groups
- · Interface groups
- VLANs
- · Physical ports or interface groups that host VLANs
- Virtual IP (VIP) ports

Starting with ONTAP 9.5, VIP LIFs are supported and are hosted on VIP ports.

While configuring SAN protocols such as FC on a LIF, it will be associated with a WWPN.

ONTAP 9 SAN Administration Guide

The following figure illustrates the port hierarchy in an ONTAP system:



Displaying all LIFs in a SVM using the CLI

You can display information about all LIFs in a SVM. You might want to display all LIFs before you determine which LIFs might be busy and should be moved.

About this task

The operational status of a LIF is determined by whether it has been configured on a particular port and is capable of serving data. When a SVM is stopped, the associated Data LIFs and SVM Management LIFs can no longer serve data. The operational status of these LIFs changes to down.

Steps

1. To display information about all LIFs in a SVM, enter the following command: network interface show -vserver vserver name

The command displays the following information:

- Node or SVM associated with the LIF
- LIF name
- Administrative and operational status
- IP address
- Netmask
- Node and port on which the LIF is configured

A home server can be either a node or a SVM.

If data for a field is not available (for instance, the operational duplex and speed for an inactive port), the field is listed as undef.



You can get all available information by specifying the -instance parameter.

The following example displays general information about all LIFs in a SVM:

		Logical	Status	Network	Current
Current	Is				
Vserver	<u>-</u>	Interface	Admin/Oper	Address/Mask	Node
Port Home					
		_			
vs1					
		lif1	up/up	192.0.2.253/24	node-01
e0b	fal	se			
		d2	up/up	192.0.2.252/21	node-01
e0d	tru	e			
		data3	up/up	192.0.2.251/20	node-02
e0c	tru	6			

Identifying LIFs with the most connections using the CLI

You might want to migrate a data LIF if it exhibits a heavy load or throughput. To decide

whether to migrate a LIF, you can display the load on LIFs, number of connections on the port, throughput, and CPU cycles on the node.

Steps

- 1. Access the CLI as a cluster administrator.
- 2. Set the privilege level to advanced by entering the following command: set -privilege advanced For details about using the CLI in advanced mode, see the System Administration Reference.
- 3. To find the weight of each LIF, enter the following command: network interface lif-weights show A busy LIF is one that has the lowest weight.
- 4. To find the active connections on a node, enter the following command: network connections active show-clients

Note the highest client count by node.

cluster	c1::> network connection	ons active show-clients
Node	Client IP Address	Count
node1	192.0.2.253	12
	192.0.2.252	9
	192.0.2.251	12
node2	192.0.2.250	12
	192.0.2.252	9
	192.0.2.253	9
node3	customer.example.com	2
	customer.example.net	2
	customer.example.org	2

5. To find the active connections by LIF on a node and SVM, enter the following command: network connections active show-lifs

Note the highest client count per LIF.

		nections active Interface Name	
node1	vs1	clus1	30
node2			
	vs2	clus1	30
node3			
	vs3	lif1	2
	vs4	clus1	30

- 6. Check the LIFs that are sharing the same home port and home node to identify the LIFs with the most connections.
- 7. To choose the best data port, enter the following: statistics show -object port

The statistics command provides throughput and bandwidth information for Ethernet ports. Each row provides a separate counter of unique information. Value is the value for the type of object since the counter was last cleared (since ONTAP was last started).

```
cluster1::> statistics show -object port
Object: port
Instance: e0a
Start-time: 10/11/2013 13:51:41
End-time: 10/11/2013 13:51:41
Node: node1
    Counter
                                                         Value
                                                            0B
    recv-data
                                                            0
    recv-packets
                                                            0
    recv-mcasts
                                                            0
    recv-errors
    recv-dropped
                                                            0
    sent-data
                                                            0B
    sent-packets
    sent-mcasts
                                                            0
    sent-errors
                                                            0
    collisions
                                                            0
```

Identifying the best node for a busy LIF using the CLI

You can display information about all the ports in a cluster. You can view information such

as the network port role (cluster, data, or node-management), link status, maximum transmission unit (MTU), speed setting and operational status, and the port's interface group, if applicable.

Steps

1. To display port information, enter the following command: network port show

The following example displays information about network ports that have a data role and are up in the cluster:

cluster1::> network port show -role data -link up						
				Auto-Negot	Duplex	Speed (Mbps)
Node Port	Role	Link	MTU	Admin/Oper	Admin/Oper	Admin/Oper
node1						
e0M	data	up	1500	true/true	full/full	auto/100
e0b	data	up	1500	true/true	full/full	auto/1000
node2						
e0b	data	up	1500	true/true	full/full	auto/1000
						J

2. Check for destination ports that are in the same network as the source home port and home node.

For example, the destination home port and home node should be on the same VLAN where applicable.

3. To identify the least busy port, choose a data port that has the least number of connections.

Identifying the best node for a busy LIF using OnCommand System Manager

You can display information about all the ports in a cluster. You can view information such as the network port role (cluster, data, or node-management), link status, maximum transmission unit (MTU), speed setting and operational status, and the port's interface group, if applicable.

Steps

- 1. Open OnCommand System Manager.
- 2. From the **Home** tab, double-click the storage system.
- 3. In the navigation pane, expand the **Nodes** hierarchy.
- 4. To find the active connections on a node, in the navigation pane, select the icon for a node.
- 5. Click the name link of a node and then click **Configuration** > **Ports/Adapters**.
- 6. Note the highest client count by node.

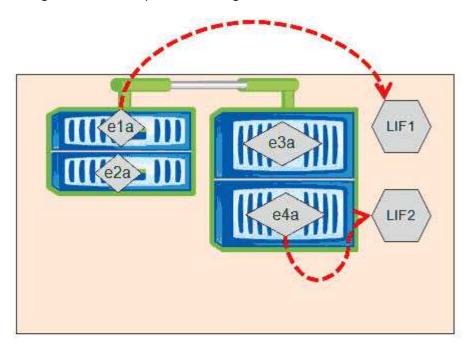
Changing home port and nodes for a LIF using OnCommand System Manager

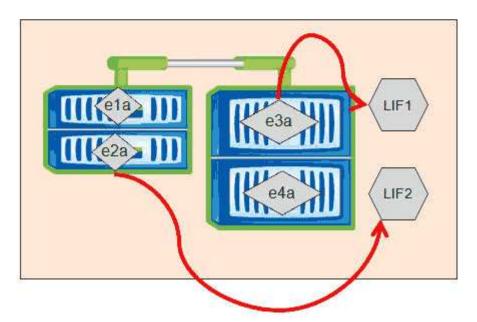
You can change the preferred location of a LIF by modifying its home port and home

node. This is a more lasting configuration than migrating a LIF, which is typically used to temporarily relocate a LIF to a different node during scheduled maintenance.

About this task

The following image shows the original LIF home port and node and the home port and node after the change. The original LIF1 home port was changed from e1a to e3a and LIF2 was changed from e4a to e2a.

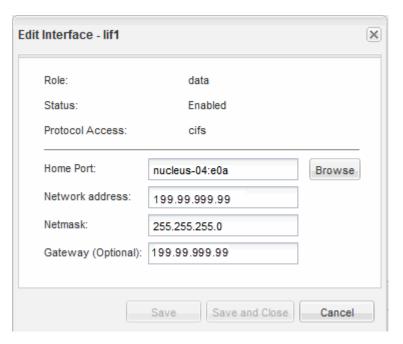




Steps

- 1. Open OnCommand System Manager.
- 2. From the **Home** tab, double-click the storage system.
- 3. In the navigation pane, expand the **SVMs** hierarchy.
- 4. In the navigation pane, select the SVMs and click **Configuration > Network Interfaces**.

- Select the LIF and click Edit.
- 6. In the Edit Interface dialog box, enter the home port and network address of the target port.



- (i)
- In ONTAP 8.2.1, the Home Port field is disabled.
- 7. Click Save and Close.

Reverting a LIF to its home port using OnCommand System Manager

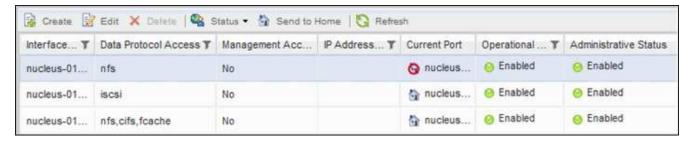
You can revert a LIF from its current port to its home port after it fails over or is migrated to a different port either manually or automatically. You can do this using OnCommand System Manager.

About this task

When creating a LIF, the administrator specifies a home port and home node to use as the preferred location of the LIF. If the home node is unavailable or the home port experiences a physical link outage, the LIF is automatically migrated to a new location. The new location is reported, in OnCommand System Manager for example, as the current port for the LIF. Unless the automatic revert option is enabled, the LIF will remain at this new location until it is reverted.

Steps

- 1. Open OnCommand System Manager.
- 2. From the **Home** tab, double-click the storage system.
- 3. In the navigation pane, expand the **Storage Virtual Machines** hierarchy.
- 4. In the navigation pane, select the SVM and click **Configuration > Network Interfaces**.
- 5. Look for data LIFs that display a house icon with a red cross mark, in the **Current Port** column, as in the following image.



6. Select the LIF and click Send to Home.

This option is enabled only when the selected interface is hosted on a non-home port and when the home port is available.

How storage QoS can control workload throughput

You can create or edit a Quality of Service (QoS) policy group to control the I/O per second (IOPS) or throughput (MBps) limit for the workloads it contains. If the workloads are in a policy group with no set limit, such as the default policy group, or the set limit does not meet your needs, you can increase the limit or move the workloads to a new or existing policy group that has the desired limit.

"Standard" QoS policy groups can be assigned to individual workloads; for example, a single volume or LUN. In this case the workload can use the full throughput limit. QoS policy groups also can be assigned to multiple workloads; in which case the throughput limit is shared among the workloads. For example, a QoS limit of 9,000 IOPS assigned to three workloads would restrict the combined IOPS from exceeding 9,000 IOPS.

"Adaptive" QoS policy groups can also be assigned to individual workloads or multiple workloads. However, even when assigned to multiple workloads, each workload gets the full throughput limit instead of sharing the throughput value with other workloads. Additionally, adaptive QoS policies automatically adjust the throughput setting based on the volume size, per workload, thereby maintaining the ratio of IOPS to terabytes as the size of the volume changes. For example, if the peak is set to 5,000 IOPS/TB in the adaptive QoS policy, a 10 TB volume will have a throughput ceiling of 50,000 IOPS. If the volume is resized later to 20 TB, adaptive QoS adjusts the ceiling to 100,000 IOPS.

Starting with ONTAP 9.5 you can include the block size when defining an adaptive QoS policy. This effectively converts the policy from an IOPS/TB threshold to a MBps threshold for cases when workloads are using very large block sizes and ultimately using a large percentage of throughput.

For shared group QoS policies, when the IOPS or MBps of all workloads in a policy group exceeds the set limit, the policy group throttles the workloads to restrict their activity, which can decrease the performance of all workloads in the policy group. If a Dynamic performance event is generated by policy group throttling, the event description displays the name of the policy group involved.

In the Performance/Volumes inventory page, you can sort the affected volumes by IOPS and MBps to see which workloads have the highest usage that might have contributed to the event. In the Performance/Volumes Explorer page, you can select other volumes or LUNs to compare to the affected workload IOPS or MBps throughput usage.

By assigning the workloads that are overusing the node resources to a more restrictive policy group setting, the policy group throttles the workloads to restrict their activity, which can reduce the use of the resources on that node. However, if you want the workload to be able to use more of the node resources, you can increase the value of the policy group.

You can use System Manager or the ONTAP commands to manage policy groups, including the following tasks:

- · Creating a policy group
- Adding or removing workloads in a policy group
- · Moving a workload between policy groups
- Changing the throughput limit of a policy group
- · Moving a workload to a different aggregate and/or node

Run storage efficiency operations at less busy times

You can modify the policy or schedule that handles storage efficiency operations to run when the impacted volume workloads are less busy.

Storage efficiency operations can use a high amount of cluster CPU resources and become a bully to the volumes on which the operations are being run. If the victim volumes have high activity at the same time when the storage efficiency operations are run, their latency can increase and trigger an event.

On the Event details page, the System Diagnosis section displays workloads in the QoS policy group by peak deviation in activity to identify the bully workloads. If you see "storage efficiency" displayed near the top of the table, these operations are bullying the victim workloads. By modifying the efficiency policy or schedule to run when these workloads are less busy, you can prevent the storage efficiency operations from causing contention on a cluster.

You can use OnCommand System Manager to manage efficiency policies. You can use the ONTAP commands to manage efficiency policies and schedules.

What storage efficiency is

Storage efficiency enables you to store the maximum amount of data for the lowest cost and accommodates rapid data growth while consuming less space. NetApp strategy for storage efficiency is based on the built-in foundation of storage virtualization and unified storage provided by its core ONTAP operating system and Write Anywhere File Layout (WAFL) file system.

Storage efficiency includes using technologies such as thin provisioning, Snapshot copy, deduplication, data compression, FlexClone, thin replication with SnapVault and volume SnapMirror, RAID-DP, Flash Cache, Flash Pool aggregate, and FabricPool-enabled aggregates which help to increase storage utilization and decrease storage costs.

The unified storage architecture allows you to efficiently consolidate a storage area network (SAN), network-attached storage (NAS), and secondary storage on a single platform.

High-density disk drives, such as serial advanced technology attachment (SATA) drives configured within Flash Pool aggregate or with Flash Cache and RAID-DP technology, increase efficiency without affecting performance and resiliency.

A FabricPool-enabled aggregate includes an all SSD aggregate as the performance tier and an object store that you specify as the cloud tier. Configuring FabricPool helps you manage which storage tier (the local performance tier or the cloud tier) data should be stored based on whether the data is frequently accessed.

Technologies such as thin provisioning, Snapshot copy, deduplication, data compression, thin replication with SnapVault and volume SnapMirror, and FlexClone offer better savings. You can use these technologies individually or together to achieve maximum storage efficiency.

Add disks and reallocate data

You can add disks to an aggregate to increase the storage capacity and the performance of that aggregate. After adding the disks, you will see an improvement in read performance only after reallocating the data across the disks you added.

You can use these instructions when Unified Manager has received aggregate events triggered by dynamic thresholds or by system-defined performance thresholds:

• When you have received a dynamic threshold event, on the Event details page, the cluster component icon that represents the aggregate in contention is highlighted red.

Beneath the icon, in parentheses, is the name of the aggregate, which identifies the aggregate to which you can add disks.

• When you have received a system-defined threshold event, on the Event details page, the event description text lists the name of the aggregate that is having the problem.

You can add disks and reallocate data on this aggregate.

The disks you add to the aggregate must already exist in the cluster. If the cluster does not have extra disks available, you might need to contact your administrator or purchase more disks. You can use OnCommand System Manager or the ONTAP commands to add disks to an aggregate.



You should reallocate data when using HDD and Flash Pool aggregates only. Do not reallocate data on SSD or FabricPool aggregates.

How enabling Flash Cache on a node can improve workload performance

You can improve workload performance by enabling Flash Cache™ intelligent data caching on each node in the cluster.

A Flash Cache module, or Performance Acceleration Module PCIe-based memory module, optimizes the performance of random read-intensive workloads by functioning as an intelligent external read cache. This hardware works in tandem with the WAFL External Cache software component of ONTAP.

In Unified Manager, on the Event details page, the cluster component icon that represents the aggregate in contention is highlighted red. Beneath the icon, in parentheses, is the name of the aggregate, which identifies the aggregate. You can enable Flash Cache on the node on which the aggregate resides.

You can use OnCommand System Manager or the ONTAP commands to see whether Flash Cache is installed or enabled, and to enable it if not already enabled. The following command indicates whether Flash Cache is enabled on a specific node: cluster::> run local options flexscale.enable

For more information about Flash Cache and the requirements for using it, see the following technical report:

Technical Report 3832: Flash Cache Best Practices Guide

How enabling Flash Pool on a storage aggregate can improve workload performance

You can improve workload performance by enabling the Flash Pool feature on an aggregate. A Flash Pool is an aggregate that incorporates both HDDs and SSDs. The HDDs are used for primary storage and the SSDs provide a high-performance read and write cache to boost aggregate performance.

In Unified Manager, the Event details page displays the name of the aggregate in contention. You can use OnCommand System Manager or the ONTAP commands to see whether Flash Pool is enabled for an aggregate. If you have SSDs installed, you can use the command-line interface to enable it. If you have SSDs installed, you can run the following command on the aggregate to see whether Flash Pool is enabled: cluster::> storage aggregate show -aggregate aggr_name -field hybrid-enabled

In this command, aggr name is the name of the aggregate, such as the aggregate in contention.

For more information about Flash Pool and the requirements for using it, see the *Clustered Data ONTAP Physical Storage Management Guide*.

MetroCluster configuration health check

You can use Unified Manager to review the health of the clusters in a MetroCluster configuration. The health status and events help you determine whether there are hardware or software issues that might be impacting the performance of your workloads.

If you configure Unified Manager to send email alerts, you can check your email for any health issues on the local or remote cluster that might have contributed to a performance event. In the Unified Manager GUI, you can select **Events** to see a list current events and then use the filters to display MetroCluster configuration events only.

MetroCluster configuration verification

You can prevent performance issues for mirrored workloads in a MetroCluster configuration by ensuring that the MetroCluster configuration is set up correctly. You can also improve workload performance by changing the configuration or upgrading software or hardware components.

The MetroCluster Installation and Configuration Guide provides instructions for setting up the clusters in the MetroCluster configuration, including the Fibre Channel (FC) switches, cables, and inter-switch links (ISLs). It also helps you configure the MetroCluster software so that the local and remote clusters can communicate with mirror volume data.

You can compare your MetroCluster configuration to the requirements in the *MetroCluster Installation and Configuration Guide* to determine whether changing or upgrading components in your MetroCluster configuration might improve workload performance. This comparison can help you answer the following questions:

- Are the controllers appropriate for your workloads?
- Do you need to upgrade your ISL bundles to a larger bandwidth to handle more throughput?
- · Can you adjust the buffer-to-buffer credits (BBC) on your switches to increase the bandwidth?

• If your workloads have high write throughput to solid state drive (SSD) storage, do you need to upgrade your FC-to-SAS bridges to accommodate the throughput?

For information about replacing or upgrading MetroCluster components, see the MetroCluster Service Guide.

Moving workloads to a different aggregate

You can use Unified Manager to help identify an aggregate that is less busy than the aggregate where your workloads currently reside, and then you can move selected volumes or LUNs to that aggregate. Moving high performing workloads to a less busy aggregate, or an aggregate with flash storage enabled, allows the workload to perform more efficiently.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have recorded the name of the aggregate that is currently having a performance issue.
- You must have recorded the date and time at which the aggregate received the event.
- You must have recorded the event ID, for example," p-sdt-clus1-ag-2542".
- Unified Manager must have collected and analyzed a month or more of performance data.

About this task

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized aggregate:

- · The aggregates on the same cluster that are less utilized
- · The highest-performing volumes on the current aggregate

Steps

- 1. Identify the aggregate in the cluster that is the least utilized:
 - a. From the **Event** details page, click the name of the cluster on which the aggregate resides.

The cluster details are displayed in the Performance/Cluster Landing page.

b. On the **Summary** page, click **Aggregates** from the **Managed Objects** pane.

The list of aggregates on this cluster are displayed.

c. Click the **Utilization** column to sort the aggregates by least utilized.

You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

- d. Write down the name of the aggregate to which you want to move the workloads.
- 2. Identify the high-performing volumes from the aggregate that received the event:
 - a. Click the aggregate that is having the performance issue.

The aggregate details are displayed in the Performance/Aggregate Explorer page.

b. From the Time Range selector, select Last 30 Days, and then click Apply Range.

This enables you to view a longer performance history period than the default 72 hours. You want to move a volume that is using a lot of resources on a consistent basis, not just over the past 72 hours.

c. From the View and Compare control, select Volumes on this Aggregate.

A list of FlexVol volumes and FlexGroup constituent volumes on this aggregate are displayed.

- d. Sort the volumes by highest MBps, and then by highest IOPS, to see the highest performing volumes.
- e. Write down the names of the volumes that you want to move to a different aggregate.
- 3. Move the high-performing volumes to the aggregate you identified as having low utilization.

You can perform the move operation by using OnCommand System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After you finish

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Moving workloads to a different node

You can use Unified Manager to help identify an aggregate on a different node that is less busy than the node on which your workloads are currently running, and then you can move selected volumes to that aggregate. Moving high-performing workloads to an aggregate on a less busy node allows the workloads on both nodes to perform more efficiently.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have recorded the name of the node that is currently having a performance issue.
- You must have recorded the date and time at which the node received the performance event.
- You must have recorded the event ID—for example, "p-sdt-clus1-nod-6982".
- Unified Manager must have collected and analyzed performance data for a month or longer.

About this task

This procedure help you to identify the following resources so that you can move high-performing workloads to a lower utilized node:

- The nodes on the same cluster that have the greatest free performance capacity
- The aggregates on the new node that have the greatest free performance capacity
- The highest-performing volumes on the current node

Steps

1. Identify a node in the cluster that has the greatest free performance capacity:

a. On the **Event Details** page, click the name of the cluster on which the node resides.

The cluster details are displayed in the Performance/Cluster Landing page.

b. On the Summary tab, click Nodes from the Managed Objects pane.

The list of nodes on this cluster are displayed.

c. Click the Performance Capacity Used column to sort the nodes by least percentage used.

This provides a list of potential nodes to which you might want to move workloads.

- d. Write down the name of the node to which you want to move the workloads.
- 2. Identify an aggregate on the new node that is the least utilized:
 - a. In the left navigation pane, click **Performance > Aggregates**.

The Performance/Aggregates page is displayed.

b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

The Performance/Aggregates page is redisplayed with the list of aggregates that are available on this node.

c. Click the Performance Capacity Used column to sort the aggregates by least used.

This provides a list of potential aggregates to which you might want to move workloads.

- d. Write down the name of the aggregate to which you want to move the workloads.
- 3. Identify the high-performing workloads from the node that received the event:
 - a. Return to the **Event Details** page for the event.
 - b. In the Affected Volumes field, click the link for the number of volumes.

The Performance/Volumes page is displayed with a filtered list of the volumes on that node.

c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you may want to move.

- d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.
- 4. Move the volumes to the aggregates that you identified as having greatest free performance capacity on the new node.

You can perform the move operation by using OnCommand System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After you finish

After a few days, you can check whether you are receiving the same type of events from this node or aggregate.

Moving workloads to an aggregate on a different node

You can use Unified Manager to help identify an aggregate on a different node that is less busy than the node where your workloads are currently running, and then you can move selected volumes to that aggregate. Moving high-performing workloads to an aggregate on a less busy node allows workloads on both nodes to perform more efficiently.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have recorded the name of the node that is currently having a performance issue.
- You must have recorded the date and time at which the node received the performance event.
- You must have recorded the event ID, for example, "p-sdt-clus1-nod-6982".
- Unified Manager must have collected and analyzed a month or more of performance data.

About this task

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized node:

- · The nodes on the same cluster that are less utilized
- · The aggregates on the new node that are the least utilized
- The highest-performing volumes on the current node

Steps

- 1. Identify a node in the cluster that is the least utilized:
 - a. From the **Event** details page, click the name of the cluster on which the node resides.

The cluster details are displayed in the Performance/Cluster Landing page.

b. On the **Summary** page, click **Nodes** from the **Managed Objects** pane.

The list of nodes on this cluster are displayed.

c. Click the **Utilization** column to sort the nodes by least utilized.

You can also identify those nodes that have the greatest **Free Capacity**. This provides a list of potential nodes to which you might want to move workloads.

- d. Write down the name of the node to which you want to move the workloads.
- 2. Identify an aggregate on the new node that is the least utilized:
 - a. In the left navigation pane, click **Performance > Aggregates**.

The Performance/Aggregates page is displayed.

b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

The Performance/Aggregates is redisplayed with the list of aggregates that are available on this node.

c. Click the **Utilization** column to sort the aggregates by least utilized.

You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

- d. Write down the name of the aggregate to which you want to move the workloads.
- Identify the high-performing workloads from the node that received the event:
 - a. Return to the **Event** details page for the event.
 - b. In the Affected Volumes field, click the link for the number of volumes.

The Performance/Volumes page is displayed with a filtered list of the volumes on that node.

c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you may want to move.

- d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.
- 4. Move the volumes to the aggregates you identified as having low utilization on the new node.

You can perform the move operation by using OnCommand System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After you finish

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Moving workloads to a node in a different HA pair

You can use Unified Manager to help identify an aggregate on a node in a different high-availability (HA) pair that has more free performance capacity than the HA pair where your workloads are currently running. Then you can move selected volumes to aggregates on the new HA pair.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- Your cluster must consist of a minimum of two HA pairs

You cannot use this remediation process if you have only one HA pair in your cluster.

- You must have recorded the names of the two nodes in the HA pair that are currently having a performance issue.
- You must have recorded the date and time at which the nodes received the performance event.
- You must have recorded the event ID—for example, "p-sdt-clus1-nod-6982".
- Unified Manager must have collected and analyzed performance data for a month or longer.

About this task

Moving high-performing workloads to an aggregate on a node with more free performance capacity allows workloads on both nodes to perform more efficiently. This procedure help you to identify the following resources so that you can move high-performing workloads to a node that has more free performance capacity on a different HA pair:

- The nodes in a different HA pair on the same cluster that have the greatest free performance capacity
- The aggregates on the new nodes that have the greatest free performance capacity
- The highest-performing volumes on the current nodes

Steps

- 1. Identify the nodes that are part of a different HA pair on the same cluster:
 - a. On the **Event Details** page, click the name of the cluster on which the nodes reside.

The cluster details are displayed in the Performance/Cluster Landing page.

b. On the **Summary** page, click **Nodes** from the **Managed Objects** pane.

The list of nodes on this cluster is displayed in the Performance/Nodes page.

- c. Write down the names of the nodes that are in different HA pairs from the HA pair that is currently having a performance issue.
- Identify a node in the new HA pair that has the greatest free performance capacity:
 - a. On the Performance/Nodes page, click the Performance Capacity Used column to sort the nodes by least percentage used.

This provides a list of potential nodes to which you might want to move workloads.

- b. Write down the name of the node on a different HA pair to which you want to move the workloads.
- Identify an aggregate on the new node that has the greatest free performance capacity:
 - a. On the Performance/Nodes page, click the node.

The node details are displayed in the Performance/Node Explorer page.

b. In the View and Compare menu, select Aggregates on this Node.

The aggregates on this node are displayed in the grid.

c. Click the Performance Capacity Used column to sort the aggregates by least used.

This provides a list of potential aggregates to which you might want to move workloads.

- d. Write down the name of the aggregate to which you want to move the workloads.
- 4. Identify the high-performing workloads from the nodes that received the event:
 - a. Return to the **Event** details page for the event.
 - b. In the Affected Volumes field, click the link for the number of volumes for the first node.

The Performance/Volumes page is displayed with a filtered list of the volumes on that node.

c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you might want to move.

- d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.
- e. Perform steps 4c and 4d for the second node that was part of this event to identify possible volumes that you want to move from that node as well.
- 5. Move the volumes to the aggregates that you identified as having greatest free performance capacity on the new node.

You can perform the move operation by using OnCommand System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After you finish

After a few days, you can check whether you are receiving the same type of events from this node or aggregate.

Moving workloads to another node in a different HA pair

You can use Unified Manager to help identify an aggregate on a node in a different HA pair that is less busy than the HA pair where your workloads are currently running. Then you can move selected volumes to aggregates on the new HA pair. Moving high-performing workloads to an aggregate on a less busy node allows workloads on both nodes to perform more efficiently.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- Your cluster must consist of a minimum of two HA pairs; you cannot use this remediation process if you
 have only one HA pair in your cluster.
- You must have recorded the names of the two nodes in the HA pair that are currently having the performance issue.
- You must have recorded the date and time at which the nodes received the performance event.
- You must have recorded the event ID, for example, "p-sdt-clus1-nod-6982".
- Unified Manager must have collected and analyzed a month or more of performance data.

About this task

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized node on a different HA pair:

- The nodes in a different HA pair on the same cluster that are less utilized
- The aggregates on the new nodes that are the least utilized
- The highest-performing volumes on the current nodes

Steps

- 1. Identify the nodes that are part of a different HA pair on the same cluster:
 - a. In the left navigation pane, click **Performance > Clusters**.

The Performance/Clusters page is displayed.

b. Click the number in the **Node Count** field for the current cluster.

The Performance/Nodes page is displayed.

- c. Write down the names of the nodes that are in different HA pairs from the HA pair that is currently having the performance issue.
- 2. Identify a node in the new HA pair that is the least utilized:
 - a. Click the **Utilization** column to sort the nodes by least utilized.

You can also identify those nodes that have the greatest **Free Capacity**. This provides a list of potential nodes to which you might want to move workloads.

- b. Write down the name of the node to which you want to move the workloads.
- 3. Identify an aggregate on the new node that is the least utilized:
 - a. In the left navigation pane, click **Performance > Aggregates**.

The Performance/Aggregates page is displayed.

b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

The Performance/Aggregates page is redisplayed with the list of aggregates that are available on this node.

c. Click the **Utilization** column to sort the aggregates by least utilized.

You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

- d. Write down the name of the aggregate to which you want to move the workloads.
- 4. Identify the high-performing workloads from the nodes that received the event:
 - a. Return to the **Event** details page for the event.
 - b. In the Affected Volumes field, click the link for the number of volumes for the first node.

The Performance/Volumes page is displayed with a filtered list of the volumes on that node.

c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you might want to move.

- d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.
- e. Perform steps 4c and 4d for the second node that was part of this event to identify possible volumes that you want to move from that node as well.

5. Move the volumes to the aggregates you identified as having low utilization on the new node.

You can perform the move operation by using OnCommand System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After you finish

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Use QoS policy settings to prioritize the work on this node

You can set a limit on a QoS policy group to control the I/O per second (IOPS) or MBps throughput limit for the workloads it contains. If workloads are in a policy group with no set limit, such as the default policy group, or the set limit does not meet your needs, you can increase the set limit or move the workloads to a new or existing policy group that has the desired limit.

If a performance event on a node is caused by workloads overusing the node resources, the event description on the Event details page displays a link to the list of volumes involved. In the Performance/Volumes page, you can sort the affected volumes by IOPS and MBps to see which workloads have the highest usage that might have contributed to the event.

By assigning the volumes that are overusing the node resources to a more restrictive policy group setting, the policy group throttles the workloads to restrict their activity, which can reduce the use of the resources on that node.

You can use OnCommand System Manager or the ONTAP commands to manage policy groups, including the following tasks:

- · Creating a policy group
- · Adding or removing workloads in a policy group
- · Moving a workload between policy groups
- · Changing the throughput limit of a policy group

Remove inactive volumes and LUNs

When aggregate free space has been identified as an issue, you can search for unused volumes and LUNs and delete them from the aggregate. This can help to alleviate the low disk space issue.

If a performance event on an aggregate is caused by low disk space, there are a few ways you can determine which volumes and LUNs are no longer being used.

To identify unused volumes:

 On the Event details page, the Affected Objects Count field provides a link that displays the list of affected volumes.

Click the link to display the volumes on the Performance/Volumes page. From there you can sort the affected volumes by **IOPS** to see which volumes have not been active.

To identify unused LUNs:

- 1. From the Event details page, write down the name of the aggregate on which the event occurred.
- 2. In the left navigation pane, click **Performance** > **LUNs**.
- 3. Click **Filtering**, select **Aggregate** from the left drop-down menu, type the name of the aggregate in the text field, and then click **Apply Filter**.
- 4. Sort the resulting list of affected LUNs by IOPS to view the LUNs that are not active.

After you have identified the unused volumes and LUNs, you can use OnCommand System Manager or the ONTAP commands to delete those objects.

Add disks and perform aggregate layout reconstruction

You can add disks to an aggregate to increase the storage capacity and the performance of that aggregate. After adding the disks, you only see an improvement in performance after reconstructing the aggregate.

When you receive a system-defined threshold event on the Event details page, the event description text lists the name of the aggregate that is having the problem. You can add disks and reconstruct data on this aggregate.

The disks you add to the aggregate must already exist in the cluster. If the cluster does not have extra disks available, you might need to contact your administrator or purchase more disks. You can use OnCommand System Manager or the ONTAP commands to add disks to an aggregate.

Technical Report 3838: Storage Subsystem Configuration Guide

Managing quotas

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

What quota limits are

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

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

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

Viewing user and user group quotas

The Health/Storage Virtual Machines inventory page displays information about the user

and user group quotas that are configured on the SVM. You can view the name of the user or user group, limits set on the disks and files, used disk and file space, and email address for notification.

Before you begin

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

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the **Health/Storage Virtual Machines** inventory page, select a SVM and then click the **User and Group Quotas** tab.

Creating rules to generate email addresses

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

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have reviewed the guidelines on the Rules to Generate User and Group Quota Email Address page.

About this task

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

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

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

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

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

Steps

- 1. In the toolbar, click 💽, and then click **Quota Email** in the left Setup menu.
- 2. In the Setup options page, click Address Rules and then enter the rule based on your criteria.
- 3. 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.

4. Click Save and Close.

5. Verify that the email address you created is displayed in the **User and Group Quotas** tab of the **Health/Storage Virtual Machine** details page.

Creating an email notification format for user and user group quotas

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click Quota Email in the left Setup menu.
- 2. In the **Setup** options page, click **Notification Format** and then enter or modify the details in the **From**, **Subject**, and **Email Details** fields.
- 3. Click **Preview** to preview the email notification.
- 4. Click **Close** to close the preview window.
- 5. Modify the content of the email notification, if required.
- 6. Click Save and Close.

Editing user and group quota email addresses

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

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have reviewed the guidelines for creating rules.

About this task

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

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the **Health/Storage Virtual Machines** inventory page, select a SVM and then click the **User and Group Quotas** tab.
- Click Edit Email Address below the row of tabs.
- 4. In the **Edit Email Address** dialog box, perform the appropriate action:

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

Understanding more about quotas

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

Overview of the quota process

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

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

- Whether the hard limit will be reached.
 - If yes, the write operation fails when the hard limit is reached and the hard quota notification is sent.
- · Whether the soft limit will be breached
 - If yes, the write operation succeeds when the soft limit is breached and the soft quota notification is sent.
- · Whether a write operation will not exceed the soft limit
 - If yes, the write operation succeeds and no notification is sent.

About quotas

Quotas provide a way to restrict or track the disk space and number of files used by a user, group, or qtree. You specify quotas using the /etc/quotas file. Quotas are applied to a specific volume or qtree.

Why you use quotas

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

You specify a quota for the following reasons:

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

Description of quotas dialog boxes

You can use the appropriate option in the User and Group Quotas tab in the Health/Storage Virtual Machines inventory page to configure the format of the email notification that is sent when a quota-related issue occurs and to configure rules to specify email addresses based on the user quota.

Email Notification Format page

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

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

From

Displays the email address from which the email is sent, which you can modify. By default, this is the email address that is specified Setup/Notifications page.

Subject

Displays the subject of the notification email.

Email Details

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

Valid keywords are as follows:

\$EVENT NAME

Specifies the event name that caused the email notification.

\$QUOTA_TARGET

Specifies the gtree or volume on which the quota is applicable.

\$QUOTA USED PERCENT

Specifies the percentage of disk hard limit, disk soft limit, file hard limit, or file soft limit that is used by the user or user group.

\$QUOTA LIMIT

Specifies the disk hard limit or file hard limit that is reached by the user or user group and one of the following events is generated:

- User or Group Quota Disk Space Hard Limit Reached
- User or Group Quota Disk Space Soft Limit Reached
- User or Group Quota File Count Hard Limit Reached
- User or Group Quota File Count Soft Limit Reached
- \$QUOTA_USED

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

\$QUOTA USER

Specifies the user or user group name.

Command buttons

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

Preview

Displays a preview of the notification email.

Restore to Factory Defaults

Enables you to restore the notification format to the factory default values.

Save

Saves the changes made to the notification format.

Rules to Generate User and Group Quota Email Address page

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

Rules area

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

How you define rules

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

For example, if you want to use the email address qtee1@xyz.com to receive notifications about quota breaches for qtree1 and use the email address admin@xyz.com for all the other qtrees, the rules must be listed in the following order:

- if (\$QTREE == 'gtree1') then gtree1@xyz.com
- if (\$QTREE == *) then admin@xyz.com

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

```
if ($USER OR GROUP == *) then $USER OR GROUP@$DOMAIN
```

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

How you add comments

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

Rules syntax

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

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

Command buttons

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

Validate

Validates the syntax of the created rule. If there are errors during validation, the rule that generates the

error is displayed along with an error message.

Restore to Factory Defaults

Enables you to restore the address rules to the factory default values.

Save

Validates the syntax of the rule and saves the rule if there are no errors. If there are errors during validation, the rule that generates the error is displayed along with an error message.

Managing and monitoring clusters and cluster object health

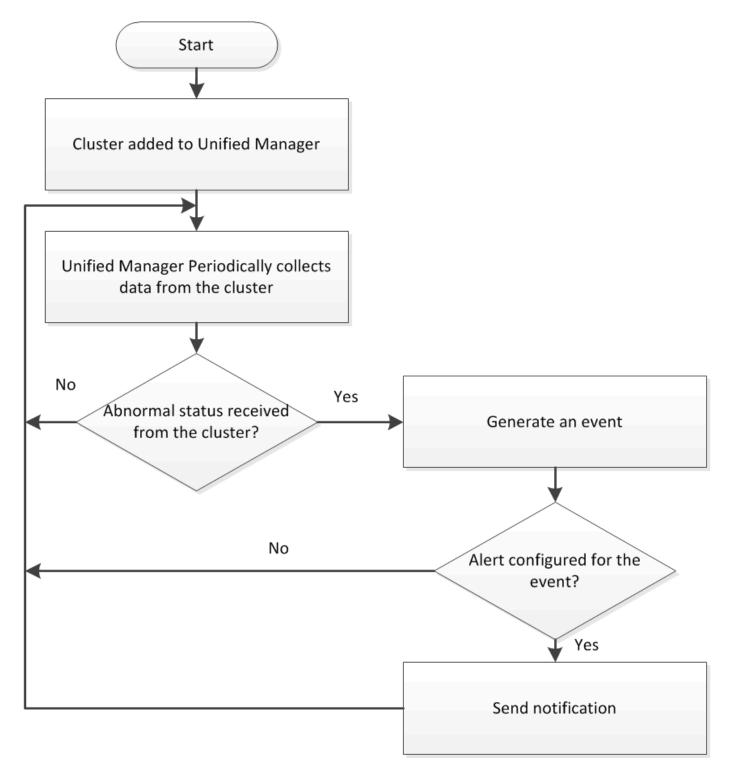
Unified Manager uses periodic API queries and a data collection engine to collect data from the clusters. By adding clusters to the Unified Manager database, you can monitor and manage these clusters for any availability and capacity risks.

Understanding cluster monitoring

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

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

The following flowchart illustrates the Unified Manager monitoring process:



Understanding node root volumes

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

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

Understanding events and thresholds for node root aggregates

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

By default, capacity and performance 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 capacity threshold values are applied to the node root aggregate.

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

Understanding quorum and epsilon

Quorum and epsilon are important measures of cluster health and function that together indicate how clusters address potential communications and connectivity challenges.

Quorum is a precondition for a fully functioning cluster. When a cluster is in quorum, a simple majority of nodes are healthy and can communicate with each other. When quorum is lost, the cluster loses the ability to accomplish normal cluster operations. Only one collection of nodes can have quorum at any one time because all of the nodes collectively share a single view of the data. Therefore, if two non-communicating nodes are permitted to modify the data in divergent ways, it is no longer possible to reconcile the data into a single data view.

Each node in the cluster participates in a voting protocol that elects one node *master*; each remaining node is a *secondary*. The master node is responsible for synchronizing information across the cluster. When quorum is formed, it is maintained by continual voting. If the master node goes offline and the cluster is still in quorum, a new master is elected by the nodes that remain online.

Because there is the possibility of a tie in a cluster that has an even number of nodes, one node has an extra fractional voting weight called *epsilon*. If the connectivity between two equal portions of a large cluster fails, the group of nodes containing epsilon maintains quorum, assuming that all of the nodes are healthy. For example, the following illustration shows a four-node cluster in which two of the nodes have failed. However, because one of the surviving nodes holds epsilon, the cluster remains in quorum even though there is not a simple majority of healthy nodes.



Epsilon is automatically assigned to the first node when the cluster is created. If the node that holds epsilon becomes unhealthy, takes over its high-availability partner, or is taken over by its high-availability partner, then epsilon is automatically reassigned to a healthy node in a different HA pair.

Taking a node offline can affect the ability of the cluster to remain in quorum. Therefore, ONTAP issues a warning message if you attempt an operation that will either take the cluster out of quorum or else put it one outage away from a loss of quorum. You can disable the quorum warning messages by using the cluster quorum-service options modify command at the advanced privilege level.

In general, assuming reliable connectivity among the nodes of the cluster, a larger cluster is more stable than a

smaller cluster. The quorum requirement of a simple majority of half the nodes plus epsilon is easier to maintain in a cluster of 24 nodes than in a cluster of two nodes.

A two-node cluster presents some unique challenges for maintaining quorum. Two-node clusters use *cluster HA*, in which neither node holds epsilon; instead, both nodes are continuously polled to ensure that if one node fails, the other has full read-write access to data, as well as access to logical interfaces and management functions.

Viewing the cluster list and details

You can use the Health/Clusters inventory page to view your inventory of clusters. The Health/Clusters Storage Summary page enables you to view summarized information about storage capacity and utilization in all clusters.

Before you begin

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

About this task

You can also view details for individual clusters such as the cluster health, capacity, configuration, LIFs, nodes, and disks in that cluster by using the Health/Cluster details page.

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

Steps

- 1. In the left navigation pane, click **Health > Clusters**.
- 2. In the **View** menu, select **Storage Summary** to view details about storage capacity and utilization in all clusters
- 3. View the complete details of the cluster in the **Health/Cluster** details page by clicking the cluster name.

Checking the health of clusters in a MetroCluster configuration

You can use Unified Manager to check the operational health of clusters, and their components, in a MetroCluster configuration. If the clusters were involved in a performance event detected by Unified Manager, the health status can help you determine whether a hardware or software issue contributed to the event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have analyzed a performance event for a MetroCluster configuration and obtained the name of the cluster involved.
- Both clusters in the MetroCluster configuration must be monitored by the same instance of Unified Manager.

Steps

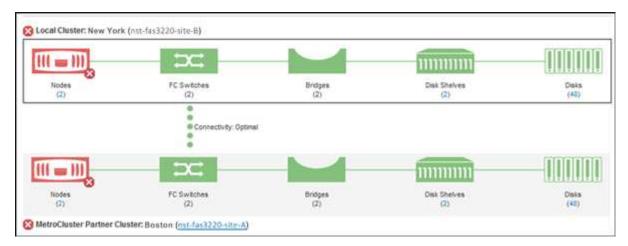
- 1. In the left navigation pane, click **Events** to display the event list.
- 2. In the filter panel, select all MetroCluster filters under the **Source Type** category.
- 3. Next to a MetroCluster event, click the name of the cluster.

The Health/Clusters inventory page is displayed with detailed information about the event.



If no MetroCluster events are displayed, you can use the Search bar to search for the name of the cluster involved in the performance event.

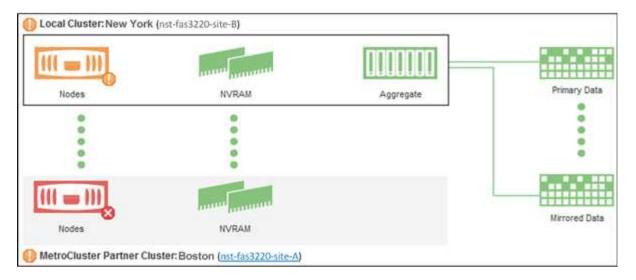
4. Select the **MetroCluster Connectivity** tab to display the health of the connection between the selected cluster and its partner cluster.



In this example, the names and the components of the local cluster and its partner cluster are displayed. A yellow or red icon indicates a health event for the highlighted component. The Connectivity icon represents the link between the clusters. You can point your mouse cursor to an icon to display event information or click the icon to display the events. A health issue on either cluster might have contributed to the performance event.

Unified Manager monitors the NVRAM component of the link between the clusters. If the FC Switches icon on the local or partner cluster or the Connectivity icon is red, a link health issue might have caused the performance event.

5. Select the MetroCluster Replication tab.



In this example, if the NVRAM icon on the local or partner cluster is yellow or red, a health issue with the NVRAM might have caused the performance event. If there are no red or yellow icons on the page, a performance issue on the partner cluster might have caused the performance event.

Viewing the node list and details

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

Before you begin

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

About this task

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

Steps

- 1. In the left navigation pane, click **Health > Nodes**.
- 2. On the Health/Nodes inventory page, click the node whose details you want to view.

The detailed information for the selected node is displayed in the Health/Cluster details page. The left pane displays the list of HA pairs. By default, the HA Details is open, which displays HA state details and events related to the selected HA pair.

3. To view other details about the node, perform the appropriate action:

To view	Click
Details about the disk shelves	Disk Shelves.

To view	Click
Port-related information	Ports.

Viewing the SVM list and details

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

Before you begin

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

About this task

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

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. Choose one of the following ways to view the SVM details:
 - To view minimal details, position the cursor over the SVM name.
 - To view the complete details, click the SVM name.

You can also view the complete details by clicking View Details in the minimal details dialog box.

3. View the objects related to the SVM by clicking **View Related** in the minimal details dialog box.

Viewing the aggregate list and details

From the Health/Aggregates inventory page, you can monitor your inventory of aggregates. The Health/Aggregates Capacity and Utilization page enables you to view information about the capacity and utilization of aggregates in all clusters.

Before you begin

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

About this task

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

Steps

1. In the left navigation pane, click **Health > Aggregates**.

- 2. Choose one of the following ways to view the aggregate details:
 - To view information about the capacity and utilization of all aggregates in all clusters, in the View menu, select **Aggregate Capacity and Utilization**.
 - To view minimal details, position the cursor over the aggregate name.
 - To view the complete details, click the aggregate name.

You can also view the complete details by clicking View Details in the minimal details dialog box.

3. View the objects related to the aggregate by clicking View Related from the minimal details dialog box.

Viewing storage pool details

You can view the details of the storage pool to monitor the storage pool health, total and available cache, and used and available allocations.

Before you begin

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

Steps

- 1. In the left navigation pane, click **Health > Aggregates**.
- 2. Click an aggregate name.

The details of the selected aggregate are displayed.

3. Click the **Disk Information** tab.

Detailed disk information is displayed.



The Cache table is displayed only when the selected aggregate is using a storage pool.

4. In the Cache table, move the pointer over the name of the required storage pool.

The details of the storage pool are displayed.

Viewing the volume list and details

From the Health/Volumes inventory page, you can monitor your inventory of volumes. The Health/Volumes Capacity and Utilization page enables you to view information about the capacity and utilization of volumes in a cluster.

Before you begin

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

About this task

You can also use the Health/Volume details page to view detailed information about volumes that are monitored, including the capacity, efficiency, configuration, and protection of the volumes. You can also view

information about the related devices and related alerts for a specific volume.

Steps

- 1. In the left navigation pane, click **Health > Volumes**.
- 2. Choose one of the following ways to view the volume details:
 - To view detailed information about the capacity and utilization of volumes in a cluster, in the View menu, select Volume Capacity and Utilization.
 - To view minimal details, position the cursor over the volume name.
 - To view the complete details, click the volume name.

You can also view the complete details by clicking View Details in the minimal details dialog box.

3. View the objects related to the volume by clicking View Related from the minimal details dialog box.

Viewing the CIFS shares

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

Before you begin

- CIFS license must be enabled on the cluster.
- · LIFs serving the CIFS shares must be configured.
- You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task



Shares in folders are not displayed in the CIFS Shares tab.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. Select the SVM for which you want to view the CIFS share details.
- 3. In the Health/Storage Virtual Machine details page, click the CIFS Shares tab.

Viewing FabricPool capacity information

You can view FabricPool capacity information for clusters, aggregates, and volumes on the Health inventory and details pages for these objects. You can also view FabricPool information in the Aggregate Capacity and Utilization report.

About this task

These pages display information such as the available capacity on the local performance tier and on the cloud tier, how much capacity is being used in both tiers, which aggregates are attached to a cloud tier, and which

volumes are implementing the FabricPool features by moving certain information to the cloud tier.

Steps

1. Perform one of the following:

To view capacity information for	Do this
Clusters	 a. On the Health/Clusters inventory page, click a cluster. b. On the Health/Cluster details page, click the Configuration tab. The display shows the names of any cloud tiers to which this cluster is connected.
Aggregates	 a. On the Health/Aggregates inventory page, click an aggregate where the Type field indicates "SSD (FabricPool)". b. On the Health/Aggregate details page, click the Capacity tab. The display shows the total capacity, plus the used and free space in the cloud tier. c. Click the Disk Information tab. The display shows the name of the cloud tier and the available space. d. Click the Configuration tab. The display shows the name of the cloud tier and other detailed information about the object store.
Volumes	 a. On the Health/Volumes inventory page, click a volume where a policy name appears in the "Tiering Policy" field. b. On the Health/Volume details page, click the Configuration tab. The display shows the name of the FabricPool tiering policy assigned to the volume.

After you finish

For more information on FabricPool aggregates, see the ONTAP 9 Disks and Aggregates Power Guide.

ONTAP 9 Disks and Aggregates Power Guide

Viewing the list of Snapshot copies

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

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- The volume containing the Snapshot copies must be online.

Steps

- 1. In the left navigation pane, click **Health > Volumes**.
- In the Health/Volumes inventory page, select the volume that contains the Snapshot copies you want to view.
- 3. In the **Health/Volume** details page, click the **Capacity** tab.
- 4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

The number of Snapshot copies is a link that displays the list of Snapshot copies.

Deleting Snapshot copies

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

The volume must be online.

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

About this task

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

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

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

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

Steps

1. In the left navigation pane, click **Health > Volumes**.

In the Health/Volumes inventory page, select the volume that contains the Snapshot copies you want to view.

The list of Snapshot copies is displayed.

- 3. In the **Health/Volume** details page, click the **Capacity** tab.
- 4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

The number of Snapshot copies is a link that displays the list of Snapshot copies.

In the Snapshot Copies view, select the Snapshot copies you want to delete, and then click Delete Selected.

Calculating reclaimable space for Snapshot copies

You can calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted.

Before you begin

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

The volume must be online.

Steps

- 1. In the left navigation pane, click **Health > Volumes**.
- 2. In the **Health/Volumes** inventory page, select the volume that contains the Snapshot copies you want to view.

The list of Snapshot copies is displayed.

- 3. In the **Health/Volume** details page, click the **Capacity** tab.
- 4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

The number of Snapshot copies is a link that displays the list of Snapshot copies.

- 5. In the **Snapshot Copies** view, select the Snapshot copies for which you want to calculate the reclaimable space.
- 6. Click Calculate.

The reclaimable space (in percentage, and KB, MB, GB, and so on) on the volume is displayed.

7. To recalculate the reclaimable space, select the required Snapshot copies and click **Recalculate**.

Description of cluster object windows and dialog boxes

You can view all your clusters and cluster objects from the respective storage object page. You can also view the details from the corresponding storage object details page.

Health/Clusters inventory page

The Health/Clusters inventory page enables you to add clusters and to view detailed information about the clusters that you are monitoring.

You must have the OnCommand Administrator or Storage Administrator role.

Command buttons

View Monitoring Status

Enables you to view the monitoring statuses of the selected clusters by navigating to the Configuration/Cluster Data Sources page.

Annotate

Enables you to annotate the selected cluster.

Refresh List

Refreshes the clusters list and the properties associated with the cluster.

Export

Enables you to export the details of all the monitored clusters to a comma-separated values (.csv) file.

Clusters table

The Clusters table displays the properties of all the discovered clusters. You can use the column filters to customize the data that is displayed:

Status

An icon that identifies the current status of the cluster. The status can be Critical (\bigotimes), Error (\bigoplus), Warning (\bigwedge), or Normal (\bigotimes).

You can position your cursor over the icon to view more information about the event or events generated for the cluster.

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

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

Cluster

The name of the cluster.

Communication Status

Whether the cluster is reachable or not.

The status is displayed as Good if the cluster is reachable. If the cluster is not reachable or if the login credentials are invalid, the status is displayed as Not Reachable.

System Health

High-level information about the status of the cluster, which is calculated based on the status of various cluster subsystems.

Possible values are OK, OK with suppressed, Degraded, and Components not reachable. These values are determined by the health monitors in ONTAP software.

Host Name or IP Address

The FQDN, short name, or the IP address of the cluster-management LIF that is used to connect to the cluster.

FQDN

The fully qualified domain name (FQDN) of the cluster.

OS Version

The ONTAP version that the cluster is running.

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

Node Count

The number of nodes that belong to the cluster.

· Last Refreshed Time

The timestamp of when the monitoring samples of the cluster were last collected.

Serial Number

The serial number of the cluster.

Contact

The contact information of the cluster.

Location

The location of the cluster.

FIPS Enabled

Whether FIPS mode is enabled on the cluster.

Filters pane

The Filters pane enables you to set filters to customize the display of information in the clusters list. You can select filters in the Status, Communication Status, System Health, and Annotation columns.



The filters specified in the Filters pane override the filters specified for the columns in the clusters list.

Health/Clusters Storage Summary page

The Health/Clusters Storage Summary page enables you to view summarized information about storage capacity and utilization in all clusters. This information helps you to understand possible capacity risks and to take appropriate action to rebalance workloads.

Use the **Export** button to export the details of all the monitored clusters to a comma-separated values (.csv) file.

Cluster

The cluster name.

HA Pair

The HA pair value obtained by forming two nodes.

Model/Family

The model or family name of the cluster.

OS Version

The version of ONTAP installed on the system.

Total Raw Capacity

Displays the total physical capacity of all disks in the array.

Unconfigured Raw Capacity

The unconfigured capacity of disks whose container type is other than aggregate, broken, spare, or shared. This capacity is always higher than the physical capacity of the disk in ONTAP. For example, consider a 2 TB disk. The physical capacity of the disk is 1.6 TB in ONTAP whereas the unconfigured raw capacity in Unified Manager is 1.8 TB.

Aggregate Total Capacity

The total size of the available aggregates for the user. This includes the Snapshot copy reserve.

Aggregate Used Capacity

The capacity already in use on aggregates. This includes the capacity consumed by volumes, LUNs, and other storage efficiency technology overheads.

Aggregate Unused Capacity

The capacity that might be available for storing additional data on the aggregate. This includes the Snapshot copy reserve.

Allocated LUN Capacity

The capacity of LUNs that are mapped.

Unallocated LUN Capacity

The capacity of all LUNs not mapped to the Host.

Volume Total Capacity

The total capacity of the volumes (used plus unused).

Volume Used Capacity

The used capacity of the volumes.

Volume Unused Capacity

The unused capacity of the volumes.

Volume Protection Capacity

The capacity of volumes that have SnapMirror and SnapVault enabled.

Cluster Licensed Cloud Tier Total

The total capacity that has been licensed in the cloud tier. This field is displayed for storage providers that require a FabricPool license, for example, Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage.

Cluster Licensed Cloud Tier Used

The space used by data in the cloud tier for storage providers that require a FabricPool license.

Cluster StorageGRID Capacity Used

The space used by data in the cloud tier for storage providers that do not require a FabricPool license, for example, StorageGRID.

Health/Cluster details page

The Health/Cluster details page provides detailed information about a selected cluster, such as health, capacity, and configuration details. You can also view information about the logical interfaces (LIFs), nodes, disks, related devices, and related alerts for the cluster.

The status next to the cluster name, for example (Good), represents the communication status; whether Unified Manager can communicate with the cluster. It does not represent the failover status or overall status of the cluster.

Command buttons

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

Switch to Performance View

Enables you to navigate to the Performance/Cluster details page.



Enables you to add the selected cluster to the Favorites dashboard.

Actions

- Add Alert: Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.
- Rediscover: Initiates a manual refresh of the cluster, which enables Unified Manager to discover recent changes to the cluster.

If Unified Manager is paired with OnCommand Workflow Automation, the rediscovery operation also reacquires cached data from WFA, if any.

After the rediscovery operation is initiated, a link to the associated job details is displayed to enable tracking of the job status.

• Annotate: Enables you to annotate the selected cluster.

View Clusters

Enables you to navigate to the Health/Clusters inventory page.

Health tab

Displays detailed information about the data availability and data capacity issues of various cluster objects such as nodes, SVMs, and aggregates. Availability issues are related to the data-serving capability of the cluster objects. Capacity issues are related to the data-storing capability of the cluster objects.

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

Availability Issues

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



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

· Capacity Issues

Graphically displays the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the cluster. For example, information is displayed about aggregates that are likely to breach the set threshold values.

Capacity tab

Displays detailed information about the capacity of the selected cluster.

Capacity

Displays the data capacity graph about the used capacity and available capacity from all allocated aggregates:

Total Capacity

Displays the total capacity of the cluster. This does not include the capacity that is assigned for parity.

Used

Displays the capacity that is used by data. This does not include the capacity that is used for parity, right-sizing, and reservation.

Available

Displays the capacity available for data.

Spares

Displays the storable capacity available for storage in all the spare disks.

Provisioned

Displays the capacity that is provisioned for all the underlying volumes.

Cloud Tier

Displays capacity details about the cloud tier for FabricPool-enabled aggregates on the cluster. A FabricPool can be either licensed or unlicensed.

Used

Displays the space used by data in configured cloud tiers.

Data graph

For an Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage, the chart displays the total data capacity that has been licensed by this cluster and the amount being used by aggregates.

For a StorageGRID, the chart displays only the total capacity being used by aggregates.

Details

Displays detailed information about the used and available capacity.

Total Capacity

Displays the total capacity of the cluster. This does not include the capacity that is assigned for parity.

Used

Displays the capacity that is used by data. This does not include the capacity that is used for parity, right-sizing, and reservation.

· Available

Displays the capacity available for data.

Provisioned

Displays the capacity that is provisioned for all the underlying volumes.

Spares

Displays the storable capacity available for storage in all the spare disks.

Cloud Tier

Displays the space used by data in configured cloud tiers. For an Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage, the total data capacity that has been licensed by this cluster is also displayed.

· Capacity Breakout by Disk Type

The Capacity Breakout by Disk Type area displays detailed information about the disk capacity of the various types of disks in the cluster. By clicking the disk type, you can view more information about the disk type from the Disks tab.

Total Usable Capacity

Displays the available capacity and spare capacity of the data disks.

• HDD

Graphically displays the used capacity and available capacity of all the HDD data disks in the cluster. The dotted line represents the spare capacity of the data disks in the HDD.

∘ Flash

SSD Data

Graphically displays the used capacity and available capacity of the SSD data disks in the cluster.

SSD Cache

Graphically displays the storable capacity of the SSD cache disks in the cluster.

SSD Spare

Graphically displays the spare capacity of the SSD, data, and cache disks in the cluster.

Unassigned Disks

Displays the number of unassigned disks in the cluster.

Aggregates with Capacity Issues list

Displays in tabular format details about the used capacity and available capacity of the aggregates that have capacity risk issues.

Status

Indicates that the aggregate has a capacity-related issue of a certain severity.

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

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

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



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

Aggregate

Displays the name of the aggregate.

Used Data Capacity

Graphically displays information about the aggregate capacity usage (in percentage).

Days to Full

Displays the estimated number of days remaining before the aggregate reaches full capacity.

Configuration tab

Displays details about the selected cluster, such as IP address, serial number, contact, and location:

Cluster Overview

Management LIF

Displays the cluster-management LIF that Unified Manager uses to connect to the cluster. The operational status of the LIF is also displayed.

Host Name or IP Address

Displays the FQDN, short name, or the IP address of the cluster-management LIF that Unified Manager uses to connect to the cluster.

FQDN

Displays the fully qualified domain name (FQDN) of the cluster.

OS Version

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

Serial Number

Displays the serial number of the cluster.

Contact

Displays details about the administrator whom you should contact in case of issues with the cluster.

Location

Displays the location of the cluster.

Remote Cluster Overview

Provides details about the remote cluster in a MetroCluster configuration. This information is displayed only for MetroCluster configurations.

Cluster

Displays the name of the remote cluster. You can click the cluster name to navigate to the details page of the cluster.

Hostname or IP Address

Displays the FQDN, short name, or IP address of the remote cluster.

Serial Number

Displays the serial number of the remote cluster.

Location

Displays the location of the remote cluster.

MetroCluster Overview

Provides details about the local cluster in a MetroCluster configuration. This information is displayed only for MetroCluster configurations.

Type

Displays whether the MetroCluster type is two-node or four-node.

Configuration

Displays the MetroCluster configuration, which can have the following values:

- Stretch Configuration with SAS cables
- Stretch Configuration with FC-SAS bridge
- Fabric Configuration with FC switches



For a four-node MetroCluster, only Fabric Configuration with FC switches is supported.

Automated Unplanned Switch Over (AUSO)

Displays whether automated unplanned switchover is enabled for the local cluster. By default, AUSO is enabled for all clusters in a two-node MetroCluster configuration in Unified Manager. You can use the command-line interface to change the AUSO setting.

Nodes

Availability

Displays the number of nodes that are up () or down () in the cluster.

OS Versions

Displays the ONTAP versions that the nodes are running as well as the number of nodes running a particular version of ONTAP. For example, 9.0 (2), 8.3 (1) specifies that two nodes are running ONTAP 9.0, and one node is running ONTAP 8.3.

Storage Virtual Machines

Availability

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

• LIFs

Availability

Displays the number of non-data LIFs that are up () or down () in the cluster.

Cluster-Management LIFs

Displays the number of cluster-management LIFs.

Node-Management LIFs

Displays the number of node-management LIFs.

Cluster LIFs

Displays the number of cluster LIFs.

Intercluster LIFs

Displays the number of intercluster LIFs.

Protocols

Data Protocols

Displays the list of licensed data protocols that are enabled for the cluster. The data protocols include iSCSI, CIFS, NFS, NVMe, and FC/FCoE.

Cloud Tiers

Lists the names of the cloud tiers to which this cluster is connected. It also lists the type (Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, Alibaba Cloud Object Storage, or StorageGRID), and the states of the cloud tiers (Available or Unavailable).

MetroCluster Connectivity tab

Displays the issues and connectivity status of the cluster components in the MetroCluster configuration. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.



The MetroCluster Connectivity tab is displayed only for clusters that are in a MetroCluster configuration.

You can navigate to the details page of a remote cluster by clicking the name of the remote cluster. You can also view the details of the components by clicking the count link of a component. For example, clicking the count link of the node in the cluster displays the node tab in the details page of the cluster. Clicking the count link of the disks in the remote cluster displays the disk tab in the details page of the remote cluster.



When managing an eight-node MetroCluster configuration, clicking the count link of the Disk Shelves component displays only the local shelves of the default HA pair. Also, there is no way to display the local shelves on the other HA pair.

You can move the pointer over the components to view the details and the connectivity status of the clusters in case of any issue and to view more information about the event or events generated for the issue.

If the status of the connectivity issue between components is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. The View Details button provides more information about the event.

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

MetroCluster Replication tab

Displays the status of the data that is being replicated. You can use the MetroCluster Replication tab to ensure data protection by synchronously mirroring the data with the already peered clusters. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.



The MetroCluster Replication tab is displayed only for clusters that are in a MetroCluster configuration.

In a MetroCluster environment, you can use this tab to verify the logical connections and peering of the local

cluster with the remote cluster. You can view the objective representation of the cluster components with their logical connections. This helps to identify the issues that might occur during mirroring of metadata and data.

In the MetroCluster Replication tab, local cluster provides the detailed graphical representation of the selected cluster and MetroCluster partner refers to the remote cluster.

LIFs tab

Displays details about all the non-data LIFs that are created on the selected cluster.

• LIF

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

Operational Status

Displays the operational status of the LIF, which can be Up (), Down (), or Unknown (). The operational status of a LIF is determined by the status of its physical ports.

Administrative Status

Displays the administrative status of the LIF, which can be Up (), Down (), or Unknown (). You can control the administrative status of a LIF when you make changes to the configuration or during maintenance. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.

IP Address

Displays the IP address of the LIF.

Role

Displays the role of the LIF. Possible roles are Cluster-Management LIFs, Node-Management LIFs, Cluster LIFs, and Intercluster LIFs.

Home Port

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

Current Port

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

Failover Policy

Displays the failover policy that is configured for the LIF.

Routing Groups

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

Failover Group

Displays the name of the failover group.

Nodes tab

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

HA Details

Provides a pictorial representation of the HA state and the health status of the nodes in the HA pair. The health status of the node is indicated by the following colors:

Green

The node is in a working condition.

Yellow

The node has taken over the partner node or the node is facing some environmental issues.

Red

The node is down.

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

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

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

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

Disk Shelves

Displays information about the disk shelves in the HA pair.

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

Shelf ID

Displays the ID of the shelf where the disk is located.

Component Status

Displays environmental details of the disk shelves, such as power supplies, fans, temperature sensors, current sensors, disk connectivity, and voltage sensors. The environmental details are displayed as icons in the following colors:

Green

The environmental components are in working properly.

Grey

No data is available for the environmental components.

Red

Some of the environmental components are down.

State

Displays the state of the disk shelf. The possible states are Offline, Online, No status, Initialization required, Missing, and Unknown.

Model

Displays the model number of the disk shelf.

Local Disk Shelf

Indicates whether the disk shelf is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

Unique ID

Displays the unique identifier of the disk shelf.

Firmware Version

Displays the firmware version of the disk shelf.

Ports

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

You can also view the events generated for the ports.

You can view the following port details:

Port ID

Displays the name of the port. For example, the port names can be e0M, e0a, and e0b.

· Role

Displays the role of the port. The possible roles are Cluster, Data, Intercluster, Node-Management, and Undefined.

Type

Displays the physical layer protocol used for the port. The possible types are Ethernet, Fibre Channel, and FCoE.

WWPN

Displays the World Wide Port Name (WWPN) of the port.

Firmware Rev

Displays the firmware revision of the FC/FCoE port.

Status

Displays the current state of the port. The possible states are Up, Down, Link Not Connected. or Unknown (?).

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

Disks tab

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

Disk Pool Summary

Displays the number of disks, which are categorized by effective types (FCAL, SAS, SATA, MSATA, SSD, Array LUN, and VMDISK), and the state of the disks. You can also view other details, such as the number of aggregate, shared disks, spare disks, broken disks, unassigned disks, and unsupported disks. If you click the effective disk type count link, disks of the selected state and effective type are displayed. For example, if you click the count link for the disk state Broken and effective type SAS, all disks with the disk state Broken and effective type SAS are displayed.

Disk

Displays the name of the disk.

RAID Groups

Displays the name of the RAID group.

Owner Node

Displays the name of the node to which the disk belongs. If the disk is unassigned, no value is displayed in this column.

State

Displays the state of the disk: Aggregate, Shared, Spare, Broken, Unassigned, Unsupported or Unknown. By default, this column is sorted to display the states in the following order: Broken, Unassigned, Unsupported, Spare, Aggregate, and Shared.

Local Disk

Displays either Yes or No to indicate whether the disk is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

Position

Displays the position of the disk based on its container type: for example, Copy, Data, or Parity. By default, this column is hidden.

Impacted Aggregates

Displays the number of aggregates that are impacted due to the failed disk. You can move the pointer over the count link to view the impacted aggregates and then click the aggregate name to view details of the aggregate. You can also click the aggregate count to view the list of impacted aggregates in the Health/Aggregates inventory page.

No value is displayed in this column for the following cases:

- For broken disks when a cluster containing such disks is added to Unified Manager
- When there are no failed disks

Storage Pool

Displays the name of the storage pool to which the SSD belongs. You can move the pointer over the storage pool name to view details of the storage pool.

Storable Capacity

Displays the disk capacity that is available for use.

Raw Capacity

Displays the capacity of the raw, unformatted disk before right-sizing and RAID configuration. By default, this column is hidden.

Type

Displays the types of disks: for example, ATA, SATA, FCAL, or VMDISK.

Effective Type

Displays the disk type assigned by ONTAP.

Certain ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and spare management. ONTAP assigns an effective disk type for each disk type.

Spare Blocks Consumed %

Displays in percentage the spare blocks that are consumed in the SSD disk. This column is blank for disks other than SSD disks.

Rated Life Used %

Displays in percentage an estimate of the SSD life used, based on the actual SSD usage and the manufacturer's prediction of SSD life. A value greater than 99 indicates that the estimated endurance has been consumed, but may not indicate SSD failure. If the value is unknown, then the disk is omitted.

Firmware

Displays the firmware version of the disk.

RPM

Displays the revolutions per minute (RPM) of the disk. By default, this column is hidden.

Model

Displays the model number of the disk. By default, this column is hidden.

Vendor

Displays the name of the disk vendor. By default, this column is hidden.

Shelf ID

Displays the ID of the shelf where the disk is located.

Bay

Displays the ID of the bay where the disk is located.

Related Annotations pane

Enables you to view the annotation details associated with the selected cluster. The details include the annotation name and the annotation values that are applied to the cluster. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

Enables you to view device details that are associated with the selected cluster.

The details include properties of the device that is connected to the cluster such as the device type, size, count, and health status. You can click on the count link for further analysis on that particular device.

You can use MetroCluster Partner pane to obtain count and also details on the remote MetroCluster partner along with its associated cluster components such as nodes, aggregates, and SVMs. The MetroCluster Partner pane is displayed only for clusters in a MetroCluster configuration.

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

MetroCluster Partner

Displays the health status of the MetroCluster partner. Using the count link, you can navigate further and obtain information about the health and capacity of the cluster components.

Nodes

Displays the number, capacity, and health status of the nodes that belong to the selected cluster. Capacity indicates the total usable capacity over available capacity.

Storage Virtual Machines

Displays the number of SVMs that belong to the selected cluster.

Aggregates

Displays the number, capacity, and the health status of the aggregates that belong to the selected cluster.

Related Groups pane

Enables you to view the list of groups that includes the selected cluster.

Related Alerts pane

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

Health/Nodes inventory page

The Health/Nodes inventory page enables you to view detailed information about the nodes in a selected cluster.

Command button

Export

Enables you to export the details of all the monitored nodes to a comma-separated values (.csv) file.

Nodes list

The Nodes list displays the properties of all the discovered nodes in a cluster. You can use the column filters to customize the data that is displayed.

Status

An icon that identifies the current status of the node. The status can be Critical (\bigotimes), Error (\bigoplus), Warning (\bigwedge), or Normal (\bigotimes).

You can position your cursor over the icon to view more information about the event or events generated for the node.

Node

The name of the node.

State

The state of the node. The state can be Up or Down.

HA State

The state of the HA pair. The state can be Error, Warning, Normal, or Not applicable.

Down Time

The time that has elapsed or the timestamp since the node is offline. If the time elapsed exceeds a week, the timestamp when the node went offline is displayed.

Cluster

The name of the cluster to which the node belongs.

Model

The model of the node.

OS version

The ONTAP software version that the node is running.

All Flash Optimized

Whether the node is optimized to support only solid-state drives (SSDs).

Serial Number

The serial number of the node.

Firmware Version

The firmware version number of the node.

Owner

The name of the node's owner.

Location

The location of the node.

Aggregate Used Capacity

The amount of space used for data in the node's aggregates.

Aggregate Total Capacity

The total space available for data in the node's aggregates.

Usable Spare Capacity

The amount of available space in the node that can be used to enhance the aggregate capacity.

Usable Raw Capacity

The amount of space that is usable in the node.

Total Raw Capacity

The capacity of every unformatted disk in the node before right-sizing and RAID configuration.

SVM Count

The number of SVMs contained by the cluster.

FC Port Count

The number of FC ports contained by the node.

FCoE Port Count

The number of FCoE ports contained by the node.

Ethernet Port Count

The number of ethernet ports contained by the node.

Flash Card Size

The size of the flash cards installed on the node.

Flash Card Count

The number of flash cards installed on the node.

· Disk Shelves Count

The number of disk shelves contained by the node.

Disk Count

The number of disks in the node.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the nodes list. You can select filters related to the Status, State, and HA State columns.



The filters that are specified in the Filters pane override the filters that are specified for the columns in the Nodes list.

Health/Aggregates inventory page

The Health/Aggregates inventory page displays information about the aggregates that are monitored, and enables you to view and modify the threshold settings.

Command buttons

Edit Thresholds

Displays the Edit Aggregate Thresholds dialog box, which enables you to edit the threshold settings for one or more aggregates.

Export

Enables you to export the details of all the monitored aggregates to a comma-separated values (.csv) file.

Aggregates list

Displays, in tabular format, the properties of all the discovered aggregates. You can use the column filters to customize the data that is displayed:

Status

The current status of the aggregate. The status can be Critical (\bigotimes), Error (\bigodot), Warning (\bigwedge), or Normal (\bigotimes).

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

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

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

Aggregate

The name of the aggregate.

You can move the pointer over an aggregate to view information such as the last generated event, node that contains the aggregate, RAID type, Snapshot reserve, Snapshot copies, and space allocated in the aggregate. You can also view the number of volume move operations that are currently in progress.

State

The current state of the aggregate:

Offline

Read or write access is not allowed.

Online

Read and write access to volumes hosted on this aggregate is allowed.

Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Creating

The aggregate is being created.

· Destroying

The aggregate is being destroyed.

Failed

The aggregate cannot be brought online.

• Frozen

The aggregate is (temporarily) not serving requests.

Inconsistent

The aggregate has been marked corrupted; contact technical support.

Iron Restricted

Diagnostic tools cannot be run on the aggregate.

Mounting

The aggregate is being mounted.

Partial

At least one disk was found for the aggregate, but two or more disks are missing.

Quiesced

The aggregate is quiesced.

Quiescing

The aggregate is being quiesced.

· Reverted

The revert operation of the aggregate is completed.

Unmounted

The aggregate is offline.

Unmounting

The aggregate is being taken offline.

Unknown

Specifies that the aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

Node

The name of the storage controller that contains the aggregate.

Mirror Status

The mirror status of the aggregate:

Mirrored

The aggregate plex data is mirrored.

Mirror degraded

The aggregate plex data cannot be mirrored.

Mirror resynchronizing

The aggregate plex data is being mirrored.

Failed

The aggregate plex data mirroring failed.

Invalid configuration

The initial state before an aggregate is created.

Uninitialized

The aggregate is being created.

Unmirrored

The aggregate is not mirrored.

CP count check in progress

The aggregate has been assimilated and Unified Manager is validating that the CP counts for the plexes is similar.

Limbo

There is an issue with the aggregate labels. The ONTAP system identifies the aggregate but cannot accurately assimilate the aggregate.

· Needs CP count check

The aggregate is assimilated but the CP counts on both plexes are not yet validated to be similar.

When an aggregate is in the mirror_resynchronizing state, then the resynchronization percentage is also shown.

In Transition

Whether the aggregate has completed transition or not.

Type

The aggregate type:

- HDD
- · Hybrid

Combines HDDs and SSDs, but Flash Pool has not been enabled.

Hybrid (Flash Pool)

Combines HDDs and SSDs, and Flash Pool has been enabled.

- · SSD
- SSD (FabricPool)

Combines SSDs and a cloud tier

VMDisk (SDS)

Virtual disks within a virtual machine

VMDisk (FabricPool)

Combines virtual disks and a cloud tier

LUN (FlexArray)

For standard disks and SSD disks, this column is blank when the monitored storage system is running a version of ONTAP earlier than 8.3.

SnapLock Type

The aggregate SnapLock Type. The possible values are Compliance, Enterprise, Non-SnapLock.

Used Data Capacity

The amount of space used for data in the aggregate.

Used Data %

The percentage of space used for data in the aggregate.

Available Data Capacity

The amount of space available for data in the aggregate.

Available Data %

The percentage of space available for data in the aggregate.

Total Data Capacity

The total data size of the aggregate.

Committed Capacity

The total space committed for all of the volumes in the aggregate.

When Autogrow is enabled on volumes that reside on the aggregate, the committed capacity is based on the maximum volume size set by autogrow, not based on the original volume size. For FabricPool aggregates, this value is relevant only to the local, or performance tier, capacity. The amount of space available in the cloud tier is not reflected in this value.

Space Savings

The storage efficiency ratio based on the total logical space that is being used to store the data and the total physical space that would be required to store the data without using ONTAP storage efficiency technologies.

This field is populated only when the monitored storage system is running ONTAP version 9.0 or greater, and only for non-root aggregates.

RAID Type

The RAID configuration type:

- RAID 0: All the RAID groups are of type RAID 0.
- RAID 4: All the RAID groups are of type RAID 4.
- · RAID-DP: All the RAID groups are of type RAID-DP.
- RAID-TEC: All the RAID groups are of type RAID-TEC.
- Mixed RAID: The aggregate contains RAID groups of different RAID types (RAID 0, RAID 4, RAID-DP, and RAID-TEC).

Cloud Tier Space Used

The amount of space being used in the cloud tier; if the aggregate is a FabricPool aggregate.

Filters pane

Enables you to set filters to customize the way information is displayed in the aggregates list. You can select filters related to the Status column.



The filters specified in the Filters pane override the filters specified for the columns in the aggregates list.

Health/Aggregates Capacity and Utilization page

The Health/Aggregates Capacity and Utilization page enables you to view information about the capacity and utilization of aggregates in all clusters. This information enables you to understand possible capacity risks and also to view the configured, used, and unused capacity of aggregates.

Use the **Export** button to export the details of all the monitored aggregates to a comma-separated values (.csv) file.

Cluster

The cluster name.

HA Pair

The HA pair value obtained by forming two nodes.

Aggregate

The aggregate name.

Total Data Capacity

The total data capacity (used plus available).

Used Data Capacity

The used data capacity.

Used Data %

The used data capacity as a percentage.

Available Data Capacity

The available data capacity.

Available Data %

The available data capacity as a percentage.

Daily Growth Rate %

The growth rate that occurs every 24 hours in the aggregate.

Days To Full

The estimated number of days remaining before the aggregate reaches full capacity.

Space Full Threshold %

The percentage at which an aggregate is considered full.

Space Nearly Full Threshold %

The percentage at which an aggregate is considered nearly full.

Growth Rate Threshold

The aggregate's growth rate that is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event.

Growth Rate Sensitivity Threshold

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

Days Until Full Threshold

The number of days remaining before the aggregate reaches full capacity.

· Snapshot Reserve Total Capacity

The total snapshot reserve capacity of the aggregate.

Snapshot Reserve Used Capacity

The amount of space used by snapshot copies from the snapshot reserve.

Snapshot Reserve Used %

The amount of space used by Snapshot copies from the snapshot reserve as a percentage.

Snapshot Reserve Available Capacity

The amount of space available for Snapshot copies.

Snapshot Reserve Available %

The amount of space available for Snapshot copies as a percentage.

Snapshot Copies Reserve Full Threshold %

The percentage at which an aggregate has consumed all its space reserved for Snapshot copies.

Overcommitted Capacity %

The aggregate overcommitment as a percentage.

Overcommitted Threshold %

The percentage at which an aggregate is considered overcommitted.

Nearly Overcommitted Threshold %

The percentage at which an aggregate is considered nearly overcommitted.

Aggregate Type

The aggregate type:

- HDD
- Hybrid

Combines HDDs and SSDs, but Flash Pool has not been enabled.

Hybrid (Flash Pool)

Combines HDDs and SSDs, and Flash Pool has been enabled.

- · SSD
- SSD (FabricPool)

Combines SSDs and a cloud tier

VMDisk (SDS)

Virtual disks within a virtual machine

VMDisk (FabricPool)

Combines virtual disks and a cloud tier

LUN (FlexArray)

For standard disks and SSD disks, this column is blank when the monitored storage system is running an ONTAP version earlier than 8.3.

RAID Type

The RAID configuration type.

· Aggregate State

The current state of the aggregate.

SnapLock Type

Whether the aggregate is a SnapLock or non-SnapLock aggregate.

Cloud Tier Space Used

The amount of data capacity that is currently being used in the cloud tier.

Cloud Tier

The name of the cloud tier object store when it was created by ONTAP.

Health/Aggregate details page

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

Command buttons



When monitoring a FabricPool-enabled aggregate, the committed and overcommitted values on this page are relevant only to the local, or performance tier, capacity. The amount of space available in the cloud tier is not reflected in the overcommitted values. Similarly, the aggregate threshold values are relevant only to the local performance tier.

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

Switch to Performance View

Enables you to navigate to the Performance/Aggregate details page.



Enables you to add the selected aggregate to the Favorites dashboard.

Actions

Add Alert

Enables you to add an alert to the selected aggregate.

· Edit Thresholds

Enables you to modify the threshold settings for the selected aggregate.

View Aggregates

Enables you to navigate to the Health/Aggregates inventory page.

Capacity tab

The Capacity tab displays detailed information about the selected aggregate, such as its capacity, thresholds, and daily growth rate.

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

Capacity

Displays the data capacity graph and the Snapshot copies graph, which display capacity details about the aggregate:

Used

Displays the space used by data in the aggregate.

Overcommitted

Indicates that the space in the aggregate is overcommitted.

Warning

Indicates that the space in the aggregate is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

Error

Indicates that the space in the aggregate is full. If this threshold is breached, the Space Full event is generated.

· Data graph

Displays the total data capacity and the used data capacity of the aggregate. If the aggregate is overcommitted, a flag is displayed with the overcommitted capacity.

Snapshot Copies graph

This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

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

Cloud Tier

Displays capacity details about the cloud tier for FabricPool-enabled aggregates. A FabricPool can be either licensed or unlicensed.

Used

Displays the space used by data in the cloud tier.

Unavailable

Displays the space in the cloud tier for an Amazon S3, Microsoft Azure Cloud FabricPool, or IBM Cloud Object Storage object that cannot be used. This space may be shared with another FabricPool-enabled aggregate.

Data graph

For an Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage, the chart displays the total data capacity that has been licensed by this cluster, the amount being used by this aggregate, and the unusable amount from other aggregates that are using the cloud tier.

For a StorageGRID, the chart displays only the total capacity being used by this aggregate.

Details

Displays detailed information about capacity.

Total Capacity

Displays the total capacity in the aggregate.

Data Capacity

Displays the amount of space used by the aggregate (used capacity) and the amount of available space in the aggregate (free capacity).

Snapshot Reserve

Displays the used and free Snapshot capacity of the aggregate.

Overcommitted Capacity

Displays the aggregate overcommitment. Aggregate overcommitment enables you to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. When thin provisioning is in use, the total size of volumes in the aggregate can exceed the total capacity of the aggregate.



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.

Cloud Tier

For an Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage, displays the total licensed capacity, the amount used by this aggregate, the amount used by other aggregates, and the free capacity for the cloud tier. For a StorageGRID, displays only the total capacity being used by this aggregate.

Total Cache Space

Displays the total space of the solid-state drives (SSDs) or allocation units that are added to a Flash Pool aggregate. If you have enabled Flash Pool for an aggregate but have not added any SSDs, then

the cache space is displayed as 0 KB.



This field is hidden if Flash Pool is disabled for an aggregate.

· Aggregate Thresholds

Displays the following aggregate capacity thresholds:

Nearly Full Threshold

Specifies the percentage at which an aggregate is nearly full.

Full Threshold

Specifies the percentage at which an aggregate is full.

Nearly Overcommitted Threshold

Specifies the percentage at which an aggregate is nearly overcommitted.

Overcommitted Threshold

Specifies the percentage at which an aggregate is overcommitted.

· Other Details: Daily Growth Rate

Displays the disk space used in the aggregate if the rate of change between the last two samples continues for 24 hours.

For example, if an aggregate uses 10 GB of disk space at 2 pm and 12 GB at 6 pm, the daily growth rate (GB) for this aggregate is 2 GB.

Volume Move

Displays the number of volume move operations that are currently in progress:

Volumes Out

Displays the number and capacity of the volumes that are being moved out of the aggregate.

You can click the link to view more details, such as the volume name, aggregate to which the volume is moved, status of the volume move operation, and the estimated end time.

Volumes In

Displays the number and remaining capacity of the volumes that are being moved into the aggregate.

You can click the link to view more details, such as the volume name, aggregate from which the volume is moved, status of the volume move operation, and the estimated end time.

Estimated used capacity after volume move

Displays the estimated amount of used space (as a percentage, and in KB, MB, GB, and so on) in the aggregate after the volume move operations are complete.

· Capacity Overview - Volumes

Displays graphs that provide information about the capacity of the volumes contained in the aggregate. The amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume is displayed. When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

You can select the graph you want to view from the drop-down lists. You can sort the data displayed in the graph to display details such as the used size, provisioned size, available capacity, fastest daily growth rate, and slowest growth rate. You can filter the data based on the storage virtual machines (SVMs) that contain the volumes in the aggregate. You can also view details for thinly provisioned volumes. You can view the details of specific points on the graph by positioning your cursor over the area of interest. By default, the graph displays the top 30 filtered volumes in the aggregate.

Disk Information tab

Displays detailed information about the disks in the selected aggregate, including the RAID type and size, and the type of disks used in the aggregate. The tab also graphically displays the RAID groups, and the types of disks used (such as SAS, ATA, FCAL, SSD, or VMDISK). You can view more information, such as the disk's bay, shelf, and rotational speed, by positioning your cursor over the parity disks and data disks.

Data

Graphically displays details about dedicated data disks, shared data disks, or both. When the data disks contain shared disks, graphical details of the shared disks are displayed. When the data disks contain dedicated disks and shared disks, graphical details of both the dedicated data disks and the shared data disks are displayed.

RAID Details

RAID details are displayed only for dedicated disks.

Type

Displays the RAID type (RAID0, RAID4, RAID-DP, or RAID-TEC).

Group Size

Displays the maximum number of disks allowed in the RAID group.

Groups

Displays the number of RAID groups in the aggregate.

Disks Used

Effective Type

Displays the types of data disks (for example, ATA, SATA, FCAL, SSD, or VMDISK) in the aggregate.

Data Disks

Displays the number and capacity of the data disks that are assigned to an aggregate. Data disk

details are not displayed when the aggregate contains only shared disks.

Parity Disks

Displays the number and capacity of the parity disks that are assigned to an aggregate. Parity disk details are not displayed when the aggregate contains only shared disks.

Shared Disks

Displays the number and capacity of the shared data disks that are assigned to an aggregate. Shared disk details are displayed only when the aggregate contains shared disks.

Spare Disks

Displays the disk effective type, number, and capacity of the spare data disks that are available for the node in the selected aggregate.



When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

SSD Cache

Provides details about dedicated cache SSD disks and shared cache SSD disks.

The following details for the dedicated cache SSD disks are displayed:

RAID Details

Type

Displays the RAID type (RAID0, RAID4, RAID-DP or RAID-TEC).

Group Size

Displays the maximum number of disks allowed in the RAID group.

Groups

Displays the number of RAID groups in the aggregate.

Disks Used

Effective Type

Indicates that the disks used for cache in the aggregate are of type SSD.

Data Disks

Displays the number and capacity of the data disks that are assigned to an aggregate for cache.

Parity Disks

Displays the number and capacity of the parity disks that are assigned to an aggregate for cache.

Spare Disks

Displays the disk effective type, number, and capacity of the spare disks that are available for the node

in the selected aggregate for cache.



When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

Provides the following details for the shared cache:

Storage Pool

Displays the name of the storage pool. You can move the pointer over the storage pool name to view the following details:

Status

Displays the status of the storage pool, which can be healthy or unhealthy.

Total Allocations

Displays the total allocation units and the size in the storage pool.

Allocation Unit Size

Displays the minimum amount of space in the storage pool that can be allocated to an aggregate.

Disks

Displays the number of disks used to create the storage pool. If the disk count in the storage pool column and the number of disks displayed in the Disk Information tab for that storage pool do not match, then it indicates that one or more disks are broken and the storage pool is unhealthy.

Used Allocation

Displays the number and size of the allocation units used by the aggregates. You can click the aggregate name to view the aggregate details.

Available Allocation

Displays the number and size of the allocation units available for the nodes. You can click the node name to view the aggregate details.

Allocated Cache

Displays the size of the allocation units used by the aggregate.

Allocation Units

Displays the number of allocation units used by the aggregate.

Disks

Displays the number of disks contained in the storage pool.

Details

Storage Pool

Displays the number of storage pools.

Total Size

Displays the total size of the storage pools.

Cloud Tier

Displays the name of the cloud tier, if you have configured a FabricPool-enabled aggregate, and shows the total licensed capacity for Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage objects.

Configuration tab

The Configuration tab displays details about the selected aggregate, such as its cluster node, block type, RAID type, RAID size, and RAID group count:

Overview

Node

Displays the name of the node that contains the selected aggregate.

Block Type

Displays the block format of the aggregate: either 32-bit or 64-bit.

· RAID Type

Displays the RAID type (RAID0, RAID4, RAID-DP, RAID-TEC or Mixed RAID).

RAID Size

Displays the size of the RAID group.

RAID Groups

Displays the number of RAID groups in the aggregate.

SnapLock Type

Displays the SnapLock Type of the aggregate.

Cloud Tier

If this is a FabricPool-enabled aggregate, the details for the object store are displayed. Some fields are different depending on the storage provider:

Name

Displays the name of the object store when it was created by ONTAP.

Object Storage Provider

Displays the name of the storage provider, for example, StorageGRID, Amazon S3, IBM Cloud Object Storage, Microsoft Azure Cloud, or Alibaba Cloud Object Storage.

Object Store Name (FQDN) or Server name

Displays the FQDN of the object store.

· Access Key or Account

Displays the access key or account for the object store.

Bucket Name or Container Name

Displays the bucket or container name of the object store.

• SSL

Displays whether SSL encryption is enabled for the object store.

History area

The History area displays graphs that provide information about the capacity of the selected aggregate. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

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

History graphs display the following information:

Aggregate Capacity Used (%)

Displays the used capacity in the aggregate and the trend in how aggregate capacity is used based on the usage history as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Capacity Used legend, the Capacity Used graph line is hidden.

Aggregate Capacity Used vs Total Capacity

Displays the trend in how aggregate capacity is used based on the usage history, as well as the used capacity and the total capacity, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Aggregate Capacity Used (%) vs Committed (%)

Displays the trend in how aggregate capacity is used based on the usage history, as well as the committed space as line graphs, as a percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Space Committed legend, the Space Committed graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity

Displays the severity of the event.

Event

Displays the event name.

Triggered Time

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp for when the event was generated is displayed.

Related Devices pane

The Related Devices pane enables you to view the cluster node, volumes, and disks that are related to the aggregate:

Node

Displays the capacity and the health status of the node that contains the aggregate. Capacity indicates the total usable capacity over available capacity.

Aggregates in the Node

Displays the number and capacity of all the aggregates in the cluster node that contains the selected aggregate. The health status of the aggregates is also displayed, based on the highest severity level. For example, if a cluster node contains ten aggregates, five of which display the Warning status and the remaining five of which display the Critical status, then the status displayed is Critical.

Volumes

Displays the number and capacity of FlexVol volumes and FlexGroup volumes in the aggregate; the number does not include FlexGroup constituents. The health status of the volumes is also displayed, based on the highest severity level.

Resource Pool

Displays the resource pools related to the aggregate.

Disks

Displays the number of disks in the selected aggregate.

Related Alerts pane

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

Health/Storage Virtual Machines inventory page

The Health/Storage Virtual Machines inventory page enables you to view detailed information about the storage virtual machines (SVMs) that you are monitoring.

Command buttons

Export

Enables you to export the details of all the monitored SVMs to a comma-separated values (.csv) file.

Annotate

Enables you to annotate the selected storage virtual machine (SVM).

SVMs list

The SVMs list displays, in tabular format, the properties of all the discovered SVMs. You can use the column filters to customize the data that is displayed:

Status

The current status of the SVM. The status can be Critical (\mathbf{X}), Error (\mathbf{P}), Warning (\mathbf{A}), or Normal (\mathbf{V}).

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

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

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

Storage Virtual Machine

The name of the SVM.

You can move the pointer over each SVM to view information such as the last generated event, cluster to which the SVM belongs, volume type of the SVM, allowed protocols, and space allocated in the SVM. You can also view the details of related objects such as the cluster to which the SVM belongs, all the SVMs that belong to the cluster, and the volumes that belong to the SVM.

State

The current administrative state of the SVM. The state can be Running, Stopped, Starting, or Stopping.

Cluster

The name of the cluster to which the SVM belongs.

Allowed Volume Type

The type of volume that can be created in the SVM. The type can be InfiniteVol, FlexVol, or FlexVol/FlexGroup.



The FlexGroup type is allowed when using ONTAP 9.1 or later.

Available Data Capacity

The available data capacity of all the volumes in the SVM.

Total Data Capacity

The total data capacity of all the volumes in the SVM.

Root Volume

The name of the root volume of the SVM.

NIS State

The state of the Network Information Service (NIS). The state can be Enabled, Disabled, or Not Configured.

NIS Domain

The NIS domain name. This column is blank when the NIS server is disabled or is not configured.

DNS State

The state of the Domain Name System (DNS). The state can be Enabled, Disabled, or Not Configured.

DNS Domain

The DNS domain name.

Name Service Switch

The information type gathered from hosts. Possible values are file, LDAP, or NIS.

LDAP Enabled

Whether the LDAP protocol is enabled or not.

Allowed Protocols

The type of protocols that can be configured on the SVM. The available protocols are FC/FCoE, iSCSI, HTTP, NDMP, NVMe, NFS, and CIFS.

Maximum Allowed Volumes

The maximum allowed volumes that can be configured on the SVM.

Volume Count

The number of volumes contained by the SVM.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the SVMs list. You can select filters related to the Status, State, and Annotation columns.



The filters specified in the Filters pane override the filters specified for the columns in the SVMs list.

Health/Storage Virtual Machine details page

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



You can monitor only data SVMs.

Command buttons

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

Switch to Performance View

Enables you to navigate to the Performance/SVM details page.

Actions

Add Alert

Enables you to add an alert to the selected SVM.

· Edit Thresholds

Enables you to edit the SVM thresholds.



This button is enabled only when on the Qtrees tab, or for an SVM with Infinite Volume.

Annotate

Enables you to annotate the selected SVM.

View Storage Virtual Machines

Enables you to navigate to the Health/Storage Virtual Machines inventory page.

Health tab

The Health tab displays detailed information about data availability, data capacity, and protection issues of various objects such as volumes, aggregates, NAS LIFs, SAN LIFs, LUNs, protocols, services, NFS exports, and CIFS shares.

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

· Availability Issues

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

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

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

Capacity Issues

Displays, as a graph, the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the SVM. For example, information is displayed about aggregates that are likely to breach the set threshold values.

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

Protection Issues

Provides a quick overview of SVM protection-related health by displaying, as a graph, the total number of relationships, including relationships that have protection issues and relationships that do not have any protection-related issues. When unprotected volumes exist, clicking on the link takes you to the Health/Volumes inventory page where you can view a filtered list of the unprotected volumes on the SVM. The colors in the graph represent the different severity levels of the issues. Clicking a graph takes you to the Protection/Volume Relationships page, where you can view a filtered list of protection relationship details. The information below the graph provides details about protection issues that can impact or have already impacted the protection of data in the SVM. For example, information is displayed about volumes that have a Snapshot copy reserve that is almost full or about SnapMirror relationship lag issues.

If the selected SVM is a repository SVM, the Protection area does not display.

Capacity tab

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

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

Capacity

The Capacity area displays details about the used and available capacity allocated from all volumes:

Total Capacity

Displays the total capacity (in MB, GB, and so on) of the SVM.

Used

Displays the space used by data in the volumes that belong to the SVM.

Guaranteed Available

Displays the guaranteed available space for data that is available for volumes in the SVM.

Unguaranteed

Displays the available space remaining for data that is allocated for thinly provisioned volumes in the SVM.

Volumes with Capacity Issues

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

Status

Indicates that the volume has a capacity-related issue of an indicated severity.

You can move the pointer over the status to view more information about the capacity-related event or events generated for the volume.

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

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



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

Volume

Displays the name of the volume.

Used Data Capacity

Displays, as a graph, information about the volume capacity usage (in percentage).

Days to Full

Displays the estimated number of days remaining before the volume reaches full capacity.

Thin Provisioned

Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.

Aggregates

For FlexVol volumes, displays the name of the aggregate that contains the volume. For FlexGroup volumes, displays the number of aggregates that are used in the FlexGroup.

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

Capacity

Displays the following capacity-related details:

- Percentage of used and free data capacity
- Percentage of used and free Snapshot capacity
- Snapshot Overflow

Displays the data space that is consumed by the Snapshot copies.

Used

Displays the space used by data in the SVM with Infinite Volume.

Warning

Indicates that the space in the SVM with Infinite Volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

Error

Indicates that the space in the SVM with Infinite Volume if full. If this threshold is breached, the Space Full event is generated.

Other Details

Total Capacity

Displays the total capacity in the SVM with Infinite Volume.

Data Capacity

Displays used data capacity, available data capacity, and Snapshot overflow capacity details of the SVM with Infinite Volume.

Snapshot Reserve

Displays the used and free details of the Snapshot reserve.

System Capacity

Displays the used system capacity and available system capacity in the SVM with Infinite Volume.

Thresholds

Displays the nearly full and full thresholds of the SVM with Infinite Volume.

Storage Class Capacity Details

Displays information about the capacity usage in your storage classes. This information is displayed only if

you have configured storage classes for your SVM with Infinite Volume.

Storage Virtual Machine Storage Class Thresholds

Displays the following thresholds (in percentage) of your storage classes:

Nearly Full Threshold

Specifies the percentage at which a storage class in an SVM with Infinite Volume is considered to be nearly full.

· Full Threshold

Specifies the percentage at which the storage class in an SVM with Infinite Volume is considered full.

Snapshot Usage Limit

Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.

Configuration tab

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

Overview

Cluster

Displays the name of the cluster to which the SVM belongs.

Allowed Volume Type

Displays the type of volumes that can be created in the SVM. The type can be InfiniteVol, FlexVol, or FlexVol/FlexGroup.

Root Volume

Displays the name of the root volume of the SVM.

Allowed Protocols

Displays the type of protocols that can be configured on the SVM. Also, indicates if a protocol is up (), down (), or is not configured ().

Data LIFs

NAS

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

· SAN

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

FC-NVMe

Displays the number of FC-NVMe 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, FC-NVMe, and FlexCache. For Infinite Volume, the SAN protocols are not applicable.

Role

Displays the LIF role. The roles can be Data or Management.

Home Port

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

Current Port

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

Port Set

Displays the port set to which the LIF is mapped.

Failover Policy

Displays the failover policy that is configured for the LIF. For NFS, CIFS, and FlexCache LIFs, the default failover policy is Next Available. Failover policy is not applicable for FC and iSCSI LIFs.

Routing Groups

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

Failover Group

Displays the name of the failover group.

Qtrees tab

The Qtrees tab displays details about qtrees and their quotas. You can click the **Edit Thresholds** button if you want to edit the health threshold settings for qtree capacity for one or more qtrees.

Use the **Export** button to create a comma-separated values (.csv) file containing the details of all the monitored qtrees. When exporting to a CSV file you can choose to create a qtrees report for the current SVM, for all SVMs in the current cluster, or for all SVMs for all clusters in your data center. Some additional qtrees fields appear in the exported CSV file.



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 (\bigotimes), Error (\bigoplus), Warning (\triangle), or Normal (\bigotimes).

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.



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.

Cluster

Displays the name of the cluster containing the qtree. Appears only in the exported CSV file.

Storage Virtual Machine

Displays the storage virtual machine (SVM) name containing the qtree. Appears only in the exported CSV

file.

Volume

Displays the name of the volume that contains the gtree.

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 gtree.

Quota Type

Specifies if the quota is for a user, user group, or a gtree. Appears only in the exported CSV file.

User or Group

Displays the name of the user or user group. There will be multiple rows for each user and user group. When the quota type is qtree or if the quota is not set, then the column is empty. Appears only in the exported CSV file.

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 in the grid page and the field is blank in the CSV export data.

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. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then "Not applicable" is displayed in the grid page and the field is blank in the CSV export data.

File Hard Limit

Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as "Unlimited" for the following conditions: if the quota is set without a file hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

File Soft Limit

Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as "Unlimited" for the following conditions: if the quota is set without a file soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

User and Group Quotas tab

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

Edit Email Address command button

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

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

Configure Email Rules command button

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

Status

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

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.



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

· User or Group

Displays the name of the user or user group. If more than one user has the same quota, the names of the users are displayed as comma-separated values.

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

Type

Specifies if the quota is for a user or a user group.

Volume or Qtree

Displays the name of the volume or qtree on which the user or user group quota is specified.

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

Disk Used %

Displays the percentage of disk space used. The value is displayed as "Not applicable" if the quota is set without a disk hard limit.

Disk Hard Limit

Displays the maximum amount of disk space allocated for the quota. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as "Unlimited" if the quota is set without a disk hard limit.

Disk Soft Limit

Displays the amount of disk space allocated for the quota before a warning event is generated. The value is displayed as "Unlimited" if the quota is set without a disk soft limit. By default, this column is hidden.

Disk Threshold

Displays the threshold value set on the disk space. The value is displayed as "Unlimited" if the quota is set without a disk threshold limit. By default, this column is hidden.

Files Used %

Displays the percentage of files used in the qtree. The value is displayed as "Not applicable" if the quota is set without a file hard limit.

File Hard Limit

Displays the hard limit for the number of files permitted on the quota. The value is displayed as "Unlimited" if the quota is set without a file hard limit.

File Soft Limit

Displays the soft limit for the number of files permitted on the quota. The value is displayed as "Unlimited" if the quota is set without a file soft limit. By default, this column is hidden.

Email Address

Displays the email address of the user or user group to which notifications are sent when there is a breach

in the quotas.

NFS Exports tab

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

Use the **Export** button to create a comma-separated values (.csv) file containing the details of all the monitored NFS exports. When exporting to a CSV file you can choose to create an NFS exports report for the current SVM, for all SVMs in the current cluster, or for all SVMs for all clusters in your data center. Some additional export policy fields appear in the exported CSV file.

Status

Displays the current status of the NFS export. The status can be Error (1) or Normal (2).

Junction Path

Displays the path to which the volume is mounted. If an explicit NFS exports policy is applied to a qtree, the column displays the path of the volume through which the qtree can be accessed.

Junction Path Active

Displays whether the path to access the mounted volume is active or inactive.

Volume or Qtree

Displays the name of the volume or qtree to which the NFS export policy is applied. For Infinite Volumes, the name of the SVM with the Infinite Volume is displayed. If an NFS export policy is applied to a qtree in the volume, the column displays both the names of the volume and the qtree.

You can click the link to view details about the object in the respective details page. If the object is a qtree, links are displayed for both the qtree and the volume.

Cluster

Displays the name of the cluster. Appears only in the exported CSV file.

Storage Virtual Machine

Displays the name of the SVM with NFS export policies. Appears only in the exported CSV file.

Volume State

Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, or Mixed.

· Offline

Read or write access to the volume is not allowed.

· Online

Read and write access to the volume is allowed.

Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Mixed

The constituents of a FlexGroup volume are not all in the same state.

· Security Style

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

UNIX (NFS clients)

Files and directories in the volume have UNIX permissions.

Unified

Files and directories in the volume have a unified security style.

NTFS (CIFS clients)

Files and directories in the volume have Windows NTFS permissions.

Mixed

Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

UNIX Permission

Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.

Export Policy

Displays the rules that define the access permission for volumes that are exported. You can click the link to view details about the rules associated with the export policy such as the authentication protocols and the access permission.

When you generate a report for the NFS Exports page, all rules that belong to the export policy are exported to the CSV file. For example, if there are two rules in the export policy, you will see only one row in the NFS Exports grid page, but the exported data will have two rows corresponding to the two rules.

Rule Index

Displays the rules associated with the export policy such as the authentication protocols and the access permission. Appears only in the exported CSV file.

Access Protocols

Displays the protocols that are enabled for the export policy rules. Appears only in the exported CSV file.

Client Match

Displays the clients that have permission to access data on the volumes. Appears only in the exported CSV file.

· Read Only Access

Displays the authentication protocol used to read data on the volumes. Appears only in the exported CSV file.

Read Write Access

Displays the authentication protocol used to read or write data on the volumes. Appears only in the exported CSV file.

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.



Shares in folders are not displayed in the CIFS Shares tab.

View User Mapping command button

Launches the User Mapping dialog box.

You can view the details of user mapping for the SVM.

Show ACL command button

Launches the Access Control dialog box for the share.

You can view user and permission details for the selected share.

Status

Displays the current status of the share. The status can be Normal () or Error ().

Share Name

Displays the name of the CIFS share.

• Path

Displays the junction path on which the share is created.

Junction Path Active

Displays whether the path to access the share is active or inactive.

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 gtree, links are displayed for both qtree and volume.

Volume State

Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, or Mixed.

Offline

Read or write access to the volume is not allowed.

Online

Read and write access to the volume is allowed.

Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Mixed

The constituents of a FlexGroup volume are not all in the same state.

Security

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

UNIX (NFS clients)

Files and directories in the volume have UNIX permissions.

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.

Export Policy

Displays the name of the export policy applicable to the share. If an export policy is not specified for the SVM, the value is displayed as Not Enabled.

You can click the link to view details about the rules associated with the export policy, such as access protocols and permissions. The link is disabled if the export policy is disabled for the selected SVM.

NFS Equivalent

Specifies whether there is an NFS equivalent for the share.

SAN tab

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

LUNs tab

Displays details about the LUNs that belong to the selected SVM. You can view information such as the LUN name, LUN state (online or offline), the name of the file system (volume or qtree) that contains the LUN, the type of host operating system, the total data capacity and serial number of the LUN. You can also view information whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group.

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

Initiator Groups tab

Displays details about initiator groups. You can view details such as the name of the initiator group, the access state, the type of host operating system that is used by all the initiators in the group, and the supported protocol. When you click the link in the access state column, you can view the current access state of the initiator group.

Normal

The initiator group is connected to multiple access paths.

Single Path

The initiator group is connected to a single access path.

No Paths

There is no access path connected to the initiator group.

You can view whether initiator groups are mapped to all the LIFs or specific LIFs through a port set. When you click the count link in the Mapped LIFs column, either all LIFs are displayed or specific LIFs for a port set are displayed. LIFs that are mapped through the target portal are not displayed. The total number of initiators and LUNs that are mapped to an initiator group is displayed.

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

Initiators tab

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

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

Data Policy tab

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



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".



The file path must always start with a junction path.

Content Placement

Displays the corresponding storage class for the rule.

Rule Filter

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

Action buttons

Create

Opens the Create Rule dialog box, which enables you to create a new rule for your data policy.

Edit

Opens the Edit Rule dialog box, which enables you to modify rule properties such as directory paths, file types, and owners.

· Delete

Deletes the selected rule.

Move Up

Moves the selected rule up in the list. However, you cannot move the default rule up in the list.

Move Down

Moves the selected rule down the list. However, you cannot move the default rule down the list.

Activate

Activates the rules and changes made to the data policy in the SVM with Infinite Volume.

· Reset

Resets all changes made to the data policy configuration.

Import

Imports a data policy configuration from a file.

Export

Exports a data policy configuration to a file.

Related Devices area

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

• LUNs

Displays the total number of the LUNs associated with the selected qtree.

NFS exports

Displays the total number of NFS export policies associated with the selected qtree.

CIFS Shares

Displays the total number of CIFS shares associated with the selected qtree.

User and Group Quotas

Displays the total number of the user and user group quotas associated with the selected qtree. The health status of the user and user group quotas is also displayed, based on the highest severity level.

Related Annotations pane

The Related Annotations pane enables you to view the annotation details associated with the selected SVM. Details include the annotation name and the annotation values that are applied to the SVM. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

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

Cluster

Displays the health status of the cluster to which the SVM belongs.

Aggregates

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

Assigned Aggregates

Displays the number of aggregates that are assigned to an SVM. The health status of the aggregates is also displayed, based on the highest severity level.

Volumes

Displays the number and capacity of the volumes that belong to the selected SVM. The health status of the volumes is also displayed, based on the highest severity level. When there are FlexGroup volumes in the SVM, the count also includes FlexGroups; it does not include FlexGroup constituents.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected SVM.

Related Alerts pane

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

Storage Pool dialog box

The Storage Pool dialog box enables you to view the details of the dedicated cache of SSDs, also known as *storage pools*. You can monitor the storage pools and view details such as the storage pool health, total and available cache, and used and available allocations in the storage pool.

You can view the following storage pool details:

Status

Displays the status of the storage pool, which can be healthy or unhealthy.

Total Allocations

Displays the total allocation units and the size in the storage pool.

Allocation Unit Size

Displays the minimum amount of space in the storage pool that can be allocated to an aggregate.

Disks

Displays the number of disks used to create the storage pool. If the disk count in the storage pool column and the number of disks displayed in the Disk Information tab for that storage pool do not match, then it indicates that one or more disks are broken and the storage pool is unhealthy.

Cache Allocations

Used Allocations

Displays the number and size of the allocation units used by the aggregates. You can click the aggregate name to view the aggregate details.

Available Allocations

Displays the number and size of the allocation units available for the nodes. You can click the node name to view the aggregate details.

Health/Volumes inventory page

The Health/Volumes inventory page displays information about the volumes in the storage systems that are monitored and enables you to modify the volume threshold settings.

Command buttons

Edit Thresholds

Displays the Edit Thresholds dialog box, which enables you to edit the health threshold settings for one or more volumes.

Protect

Displays the following submenus:

SnapMirror

Enables you to create a SnapMirror relationship for the selected volumes.

SnapVault

Enables you to create a SnapVault relationship for the selected volumes.

Restore

Displays the Restore dialog box, which enables you to restore directories or files from one volume at a time.

This button is disabled if more than one volume is selected, or if a FlexGroup volume is selected, or if a volume configured for SnapMirror Synchronous is selected.

Annotate

Enables you to annotate the selected volume.

Export

Enables you to export the details of all the monitored volumes to a comma-separated values (.csv) file. When viewing Infinite Volumes, the Infinite Volume constituents will be exported.

Volumes Overview table

The volumes table displays the properties of all the discovered volumes. You can use the column filters to customize the data that is displayed:

Status

The current status of a volume. The status can be Critical (\bigotimes), Error (\bigodot), Warning (\bigwedge), or Normal (\bigotimes).

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

If the status of the volume is determined by a single event, you can view information such as the event

name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** link to view more information about the event.

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

Volume

The name of the volume.

You can move the pointer over a volume to view information such as the qtree quota overcommitted space, status of the last volume move operation, and space allocated in the volume. You can also view the details of related objects such as the SVM to which the volume belongs, the aggregate to which the volume belongs, and all the volumes that belong to this aggregate.

If an SVM with Infinite Volume is monitored, you can view details about the three types of constituents (data, namespace, and namespace mirror) in the SVM with Infinite Volume. The constituent details include the following information:

- Constituent name
- State of the constituent
- Name of the SVM with Infinite Volume to which the constituent belongs
- Junction path of the constituent
- · Name of the aggregate that contains the constituent
- Available, used, and total data capacity of the constituent

State

The current state of the volume:

· Offline

Read or write access to the volume is not allowed.

Online

Read and write access to the volume is allowed.

Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Mixed

The constituents of a FlexGroup volume are not all in the same state.

Junction Path

The path to which the volume is mounted.

Storage Virtual Machine

The SVM that contains the volume.

Aggregates

The name of the aggregate on which the volume resides, or the number of aggregates on which the FlexGroup volume resides.

You can click the name to display details in the Aggregate details page. For FlexGroup volumes, you can click the number to display the aggregates that are used in the FlexGroup in the Aggregates page.

Tiering Policy

The tiering policy set on the volume. The policy takes affect only when the volume is deployed on a FabricPool aggregate:

- None. The data for this volume always remains on the performance tier.
- Snapshot-Only. Only Snapshot data is moved automatically to the cloud tier. All other data remains on the performance tier.
- Backup. On data protection volumes, all transferred user data starts in the cloud tier, but later client reads can cause hot data to move to the performance tier.
- Auto. Data on this volume is moved between the performance tier and the cloud tier automatically when ONTAP determines that the data is "hot" or "cold".

SnapLock Type

The SnapLock Type of the aggregate that contains the volume. The available options are Compliance, Enterprise, Non-SnapLock.

In Transition

Whether the volume has completed transition or not.

Protection Role

The protection role of a volume:

Unprotected

A read/write volume with no outgoing or incoming SnapMirror or SnapVault relationships

Protected

A read/write volume with an outgoing SnapMirror or SnapVault relationship

Destination

A data protection (DP) volume or read/write volume with an incoming SnapMirror or SnapVault relationship

Not Applicable

A volume for which protection roles do not apply, such as a load sharing volume, data constituent, or temporary volume

You can move your pointer over the protection role for a volume to display a graphical representation of the protection topology for the selected volume. This may include the source volume, the total number of outgoing synchronous and asynchronous SnapMirror relationships, and the total number of outgoing SnapVault relationships. Blue highlighting around the volume indicates the selected volume.

Clicking View Protection Details displays the Protection tab of the Health/Volume details page.

Thin Provisioned

Whether space guarantee is set for the selected volume. Valid values are Yes and No.

Available Data Capacity

The amount of physical space currently available for data in the volume.

Available Data %

The percentage of physical space currently available for data in the volume.

Used Data Capacity

The amount of physical space used by data in the volume.

Used Data %

The percentage of physical space used by data in the volume based on the total available data capacity.

Logical Space Used %

The percentage of logical space used by data in the volume based on the total available data capacity.

Logical Space Reporting

Whether the volume has logical space reporting configured. The value can be Enabled, Disabled, or Not applicable.

Logical space indicates the real size of the data that is being stored on the volume without applying the savings from using ONTAP storage efficiency technologies.

Total Data Capacity

The total physical space available for data in the volume.

Storage Class

The storage class name. This column is displayed for Infinite Volume only.

Constituent Role

The role name of the constituent. The roles can be Namespace, Data, or Namespace Mirror. This column is displayed for Infinite Volumes only.

Move Status

The current status of the volume move operation. The status can be In Progress, Paused, Failed, or Completed.

You can move the pointer over the status to view more information about the volume move operation, such as the source, destination, operation start time, operation end time, current phase of the volume move operation that is in progress, status (in percentage), and estimated end time.

Caching Policy

The caching policy that is associated with the selected volume. The policy provides information about how the Flash Pool caching occurs for the volume.

Cache policy	Description
Auto	Read caches all the metadata blocks and randomly read user data blocks, and write caches all the randomly overwritten user data blocks.
None	Does not cache any user data or metadata blocks.
All	Read caches all the user data blocks that are read and written. The policy does not perform any write caching.
All-Random Write	 This policy is a combination of the All and No Read-Random Write policies and performs the following actions: Read caches all the user data blocks that are read and written. Write caches all the randomly overwritten user data blocks.
All Read	Read caches all the metadata, randomly read, and sequentially read user data blocks.
All Read-Random Write	 This policy is a combination of the All Read and No Read-Random Write policies and performs the following actions: Read caches all the metadata, randomly read, and sequentially read user data blocks. Write caches all the randomly overwritten user data blocks.
All Read Random Write	Read caches all the metadata, randomly read, sequentially read, and randomly written user data blocks.

Cache policy	Description
All Read Random Write-Random Write	 This policy is a combination of the All Read Random Write and No Read-Random Write policies and does the following: Read caches all the metadata, randomly read, and sequentially read, and randomly written user data blocks. Write caches all the randomly overwritten user data blocks.
Meta	Read caches only metadata blocks.
Meta-Random Write	This policy is a combination of the Meta and No Read-Random Write and does the following: Read caches only
No Read-Random Write	Write caches all the randomly overwritten user data blocks. The policy does not perform any read caching.
Random Read	Read caches all the metadata blocks and randomly read user data blocks.
Random Read-Write	Read caches all the metadata, randomly read, and randomly written user data blocks.
Random Read-Write-Random Write	This policy is a combination of the Random Read Write and No Read-Random Write policies and does the following: • Read caches all the metadata, randomly read, and randomly overwritten user data blocks. • Write caches all the randomly overwritten user data blocks.

Cache Retention Priority

The cache retention priority for the volume. A cache retention priority defines how long the blocks of a volume will be in cache state in a Flash Pool once they become cold.

Low

Cache the cold volume blocks for the lowest time

Normal

Cache the cold volume blocks for the default time

High

Cache the cold volume blocks for the highest time

Compression

Whether compression is enabled on the volume. The column displays either Enabled or Disabled.

Deduplication

Whether deduplication is enabled on the volume. The column displays either Enabled or Disabled.

Style

The style of volume; FlexVol or FlexGroup.

Type

The volume type. The volume type can be Read-write or Data-protection, Load-sharing, or Data-cache.

Cluster

The cluster that contains the destination volume. You can view more details about the cluster by clicking the cluster name.

Cluster Nodes

The name of the node to which the volume belongs, or the number of nodes on which the FlexGroup volume resides. You can view more details about the cluster node by clicking the node name.

You can click the node name to display details in the Node details page. For FlexGroup volumes, you can click the number to display the nodes that are used in the FlexGroup in the Nodes page.

Local Snapshot Policy

The local Snapshot copy policies for the volumes listed. The default policy name is Default.

Filters pane

The Filters pane enables you to set filters to customize the way information is displayed in the volumes list. You can select filters related to the Volume Status, State, and Annotation columns.



The filters specified in the Filters pane override the filters specified for the columns in the volumes list.

Health/Volumes Capacity and Utilization page

The Health/Volumes Capacity and Utilization page enables you to view information about the capacity and utilization of volumes in a cluster. This information enables you to understand possible capacity risks and to view the configured, used, and unused capacity of volumes. Also, the information helps you to make decisions about enabling spacesaving features such as deduplication and thin provisioning.

Use the **Export** button to export the details of all the monitored volumes to a comma-separated values (.csv) file.

Cluster

The cluster name.

SVM

The name of the storage virtual machine (SVM) that contains the volume.

Volume

The volume name.

Total Data Capacity

The total data capacity (used plus available) in a volume.

Used Data Capacity

The used data capacity in a volume.

Used Data %

The used data in a volume as a percentage.

Available Data Capacity

The available data capacity in a volume.

Available Data %

The available data capacity in a volume as a percentage.

Daily Growth Rate %

The growth rate that occurs every 24 hours in the volume.

Days To Full

The estimated number of days remaining before the volume reaches full capacity.

Space Full Threshold %

The percentage of space used in the volume that is considered full.

Space Nearly Full Threshold %

The percentage of space used in the volume that is considered nearly full.

Growth Rate Threshold %

The volume's growth rate that is considered to be normal before the system generates a Volume Growth Rate Abnormal event.

· Growth Rate Sensitivity Threshold

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

Days Until Full Threshold

The number of days remaining before reaching full capacity.

Snapshot Overflow %

The percentage of the data space that is consumed by Snapshot copies.

Snapshot Reserve Used Capacity

The amount of space used by Snapshot copies in the volume.

Snapshot Reserve Used %

The amount of space used by Snapshot copies in the volume as a percentage.

Snapshot Reserve Available Capacity

The amount of space available for Snapshot copies in the volume.

Snapshot Reserve Available %

The amount of space available for Snapshot copies in the volume as a percentage.

Snapshot Reserve Total Capacity

Displays the total Snapshot copy capacity in the volume.

Snapshot Copies Reserve Full Threshold %

The percentage at which the space reserved for Snapshot copies is considered full.

Snapshot Copies Count Threshold

The number of Snapshot copies on a volume that are considered to be too many.

Snapshot Copies Days Until Full Threshold

The number of days remaining before the space reserved for Snapshot copies reaches full capacity.

Number Of Inodes

The number of inodes in the volume.

Inode Utilization %

The percentage of inode space used in the volume.

Inodes Full Threshold %

The percentage at which a volume is considered to have consumed all of its inodes.

Inodes Nearly Full Threshold %

The percentage at which a volume is considered to have consumed most of its inodes.

Quota Committed Capacity

The space reserved for quotas in the volume.

Quota Overcommitted Capacity

The amount of space that can be used for quotas before the system generates the Volume Quota Overcommitted event.

Quota Overcommitted Threshold %

The percentage at which the space used for quotas on the volume is considered overcommitted.

Quota Nearly Overcommitted Threshold %

The percentage at which the space used for quotas on the volume is considered nearly overcommitted.

Snapshot Autodelete

Whether automatic deletion of Snapshot copies is enabled or disabled.

Deduplication

Whether deduplication is enabled or disabled for the volume.

Deduplication Space Savings

The amount of space saved in a volume by using deduplication.

Compression

Whether compression is enabled or disabled for the volume.

Compression Space Savings

The amount of space saved in a volume by using compression.

Caching Policy

The caching policy that is associated with the selected volume.

The policy provides information about how Flash Pool caching occurs for the volume. See the Health/Volumes inventory page for more information on caching policies.

Cache Retention Priority

The priority used for retaining cached pools.

Thin Provisioned

Whether space guarantee is set for the selected volume. Valid values are Yes and No.

Autogrow

Whether the volume automatically grows in size when it is out of space.

Space Guarantee

The storage guarantee option that is associated with the volume.

Protection Role

The protection role that is set for the volume.

State

The state of the volume that is being exported.

· SnapLock Type

Whether the volume is a SnapLock or non-SnapLock volume.

SnapLock Expiry Date

The SnapLock expiration date.

Tiering Policy

The tiering policy set for the volume. Valid when deployed on FabricPool-enabled aggregates only.

Health/Volume details page

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

You must have the OnCommand Administrator or Storage Administrator role.

Command buttons

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

Switch to Performance View

Enables you to navigate to the Performance/Volume details page.



Enables you to add the selected volume to the Favorites dashboard.

Actions

Add Alert

Enables you to add an alert to the selected volume.

Edit Thresholds

Enables you to modify the threshold settings for the selected volume.

Annotate

Enables you to annotate the selected volume.

Protect

Enables you to create either SnapMirror or SnapVault relationships for the selected volume.

Relationship

Enables you to execute the following protection relationship operations:

Edit

Launches the Edit Relationship dialog box which enables you to change existing SnapMirror policies, schedules, and maximum transfer rates for an existing protection relationship.

Abort

Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer.

Quiesce

Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress must complete before the relationship is quiesced.

Break

Breaks the relationship between the source and destination volumes and changes the destination to a read-write volume.

Remove

Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed. This operation cannot be undone.

Resume

Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used, if one exists.

Resynchronize

Enables you to resynchronize a previously broken relationship.

Initialize/Update

Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform a manual update if the relationship is already initialized.

Reverse Resync

Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by the contents on the destination, and any data that is newer than the data on the common Snapshot copy is deleted.

• Restore

Enables you to restore data from one volume to another volume.



The Restore button and the Relationship operation buttons are not available for FlexGroup volumes, or for volumes that are in synchronous protection relationships.

View Volumes

Enables you to navigate to the Health/Volumes inventory page.

Capacity tab

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

· Capacity Physical

Details the physical capacity of the volume:

Snapshot Overflow

Displays the data space that is consumed by the Snapshot copies.

Used

Displays the space used by data in the volume.

Warning

Indicates that the space in the volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

Error

Indicates that the space in the volume is full. If this threshold is breached, the Space Full event is generated.

Unusable

Indicates that the Thin-Provisioned Volume Space At Risk event is generated and that the space in the thinly provisioned volume is at risk because of aggregate capacity issues. The unusable capacity is displayed only for thinly provisioned volumes.

Data graph

Displays the total data capacity and the used data capacity of the volume.

If autogrow is enabled, the data graph also displays the space available in the aggregate. The data graph displays the effective storage space that can be used by data in the volume, which can be one of the following:

- Actual data capacity of the volume for the following conditions:
 - Autogrow is disabled.
 - Autogrow-enabled volume has reached the maximum size.
 - Autogrow-enabled thickly provisioned volume cannot grow further.
- Data capacity of the volume after considering the maximum volume size (for thinly provisioned volumes and for thickly provisioned volumes when the aggregate has space for the volume to reach maximum size)
- Data capacity of the volume after considering the next possible autogrow size (for thickly provisioned volumes that have an autogrow percentage threshold)
- 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.

Capacity Logical

Displays the logical space characteristics of the volume. The logical space indicates the real size of the data that is being stored on disk without applying the savings from using ONTAP storage efficiency technologies.

Logical Space Reporting

Displays if the volume has logical space reporting configured. The value can be Enabled, Disabled, or Not applicable. "Not applicable" is displayed for volumes on older versions of ONTAP or on volumes that do not support logical space reporting.

Used

Displays the amount of logical space that is being used by data in the volume, and the percentage of logical space used based on the total data capacity.

Available

Displays the amount of logical space that is still available for data in the volume, and the percentage of logical space available based on the total data capacity.

Logical Space Enforcement

Displays whether logical space enforcement is configured for thinly provisioned volumes. When set to Enabled, the logical used size of the volume cannot be greater than the currently set physical volume size.

Autogrow

Displays whether the volume automatically grows when it is out of space.

Space Guarantee

Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate. These blocks are then guaranteed to be available for writes to files in the volume. The space guarantee can be set to one of the following:

None

No space guarantee is configured for the volume.

File

Full size of sparsely written files (for example, LUNs) is guaranteed.

Volume

Full size of the volume is guaranteed.

Partial

The FlexCache volume reserves space based on its size. If the FlexCache volume's size is 100 MB or more, the minimum space guarantee is set to 100 MB by default. If the FlexCache volume's size is less than 100 MB, the minimum space guarantee is set to the FlexCache volume's size. If the FlexCache volume's size is grown later, the minimum space guarantee is not incremented.



The space guarantee is Partial when the volume is of type Data-Cache.

· Details (Physical)

Displays the physical characteristics of the volume.

Total Capacity

Displays the total physical capacity in the volume.

Data Capacity

Displays the amount of physical space used by the volume (used capacity) and the amount of physical space that is still available (free capacity) in the volume. These values are also displayed as a percentage of the total physical capacity.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

Snapshot Reserve

Displays the amount of space used by the Snapshot copies (used capacity) and amount of space available for Snapshot copies (free capacity) in the volume. These values are also displayed as a percentage of the total snapshot reserve.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the Snapshot copies (used capacity) and the amount of space that is available in the volume but cannot be used for making Snapshot copies (unusable capacity) because of aggregate capacity issues is displayed.

Volume Thresholds

Displays the following volume capacity thresholds:

· Nearly Full Threshold

Specifies the percentage at which a volume is nearly full.

Full Threshold

Specifies the percentage at which a volume is full.

Other Details

Autogrow Max Size

Displays the maximum size up to which the 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.

Qtree Quota Committed Capacity

Displays the space reserved in the quotas.

Qtree Quota Overcommitted Capacity

Displays the amount of space that can be used before the system generates the Volume Qtree Quota Overcommitted event.

Fractional Reserve

Controls the size of the overwrite reserve. By default, the fractional reserve is set to 100, indicating that 100 percent of the required reserved space is reserved so that the objects are fully protected for overwrites. If the fractional reserve is less than 100 percent, the reserved space for all the space-reserved files in that volume is reduced to the fractional reserve percentage.

Snapshot Daily Growth Rate

Displays the change (in percentage, or in KB, MB, GB, and so on) that occurs every 24 hours in the Snapshot copies in the selected volume.

Snapshot Days to Full

Displays the estimated number of days remaining before the space reserved for the Snapshot copies in the volume reaches the specified threshold.

The Snapshot Days to Full field displays a Not Applicable value when the growth rate of the Snapshot copies in the volume is zero or negative, or when there is insufficient data to calculate the growth rate.

Snapshot Autodelete

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

Snapshot Copies

Displays information about the Snapshot copies in the volume.

The number of Snapshot copies in the volume is displayed as a link. Clicking the link opens the

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

The Snapshot copy count is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

Volume Move

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

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

Efficiency tab

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

Deduplication

Enabled

Specifies whether deduplication is enabled or disabled on a volume.

Space Savings

Displays the amount of space saved (in percentage, or in KB, MB, GB, and so on) in a volume by using deduplication.

· Last Run

Displays the time that has elapsed since the deduplication operation was last performed. Also specifies whether the deduplication operation was successful.

If the time elapsed exceeds a week, the timestamp representing when the operation was performed is displayed.

Mode

Specifies whether the deduplication operation enabled on a volume is a manual, scheduled, or policy-based operation. If the mode is set to Scheduled, the operation schedule is displayed, and if the mode is set to a policy, the policy name is displayed.

Status

Displays the current status of the deduplication operation. The status can be Idle, Initializing, Active, Undoing, Pending, Downgrading, or Disabled.

Type

Specifies the type of deduplication operation running on the volume. If the volume is in a SnapVault relationship, the type displayed is SnapVault. For any other volume, the type is displayed as Regular.

Compression

Enabled

Specifies whether compression is enabled or disabled on a volume.

Space Savings

Displays the amount of space saved (in percentage, or in KB, MB, GB, and so on) in a volume by using compression.

Configuration tab

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

Overview

Full Name

Displays the full name of the volume.

Aggregates

Displays the name of the aggregate on which the volume resides, or the number of aggregates on which the FlexGroup volume resides.

Tiering Policy

Displays the tiering policy set for the volume; if the volume is deployed on a FabricPool-enabled aggregate. The policy can be None, Snapshot Only, Backup, or Auto.

Storage Virtual Machine

Displays the name of the storage virtual machine (SVM) that contains the volume.

Junction Path

Displays the status of the path, which can be active or inactive. The path in the SVM to which the volume is mounted is also displayed. You can click the **History** link to view the most recent five changes to the junction path.

Export policy

Displays the name of the export policy that is created for the volume. You can click the link to view details about the export policies, authentication protocols, and access enabled on the volumes that belong to the SVM.

Style

Displays the volume style. The volume style can be FlexVol or FlexGroup.

Type

Displays the type of the selected volume. The volume type can be Read-write, Load-sharing, Data-Protection, Data-cache, or Temporary.

• RAID Type

Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, or RAID-TEC.



Multiple RAID types may display for FlexGroup volumes because the constituent volumes for FlexGroups can be on aggregates of different types.

SnapLock Type

Displays the SnapLock Type of the aggregate that contains the volume.

SnapLock Expiry

Displays the expiry date of SnapLock volume.

Capacity

Thin Provisioning

Displays whether thin provisioning is configured for the volume.

Autogrow

Displays whether the flexible volume grows automatically within an aggregate.

Snapshot Autodelete

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

Quotas

Specifies whether the quotas are enabled for the volume.

Efficiency

Deduplication

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

Compression

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

Protection

Snapshot Copies

Specifies whether automatic Snapshot copies are enabled or disabled.

Protection tab

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

Summary

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

Source Volume

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

Lag Status

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

The lag status is not applicable for synchronous relationships.

Lag Duration

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

Last Successful Update

Displays the date and time of the most recent successful protection update.

The last successful update is not applicable for synchronous relationships.

Storage Service Member

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

Version Flexible Replication

Displays either Yes, Yes with backup option, or None. Yes indicates that SnapMirror replication is possible even if source and destination volumes are running different versions of ONTAP software. Yes with backup option indicates the implementation of SnapMirror protection with the ability to retain multiple versions of backup copies on the destination. None indicates that Version Flexible Replication is not enabled.

Relationship Capability

Indicates the ONTAP capabilities available to the protection relationship.

Protection Service

Displays the name of the protection service if the relationship is managed by a protection partner application.

Relationship Type

Displays any relationship type, including Asynchronous Mirror, Asynchronous Vault, StrictSync, and Sync.

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:

Aborting

SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

Checking

The destination volume is undergoing a diagnostic check and no transfer is in progress.

Finalizing

SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

Idle

Transfers are enabled and no transfer is in progress.

In-Sync

The data in the two volumes in the synchronous relationship are synchronized.

Out-of-Sync

The data in the destination volume is not synchronized with the source volume.

Preparing

SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.

Queued

SnapMirror transfers are enabled. No transfers are in progress.

Quiesced

SnapMirror transfers are disabled. No transfer is in progress.

Quiescing

A SnapMirror transfer is in progress. Additional transfers are disabled.

Transferring

SnapMirror transfers are enabled and a transfer is in progress.

Transitioning

The asynchronous transfer of data from the source to the destination volume is complete, and the transition to synchronous operation has started.

Waiting

A SnapMirror transfer has been initiated, but some associated tasks are waiting to be queued.

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 Asynchronous Mirror protection policy, and XDPDefault indicates the default Asynchronous Vault policy. StrictSync indicates the default Synchronous Strict protection policy, and Sync indicates the default Synchronous 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. 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.

Double bold lines specify an Asynchronous Mirror relationship, a single bold line specifies an Asynchronous Vault relationship, and a bold line and non-bold line specifies a Synchronous relationship. The table below indicates if the relationship is StrictSync or Sync.

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
- If the volume is a FlexGroup volume
- If the volume is in a synchronous protection relationship
- When the volume ID is unknown, for example, when you have a intercluster relationship and the
 destination cluster has not yet been discovered
 Clicking another volume in the topology selects and displays information for that volume. A question
 mark (?) in the upper-left corner of a volume indicates that either the volume is missing or that it has
 not yet been discovered. It might also indicate that the capacity information is missing. Positioning your
 cursor over the question mark displays additional information, including suggestions for remedial
 action.

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

Topology views include the following information:

Capacity

Displays the total amount of capacity used by the volume. Positioning your cursor over a volume in the topology displays the current warning and critical threshold settings for that volume in the Current Threshold Settings dialog box. You can also edit the threshold settings by clicking the **Edit Thresholds** link in the Current Threshold Settings dialog box. Clearing the **Capacity** check box hides all capacity information for all volumes in the topology.

Lag

Displays the lag duration and the lag status of the incoming protection relationships. Clearing the **Lag** check box hides all lag information for all volumes in the topology. When the **Lag** check box is dimmed, then the lag information for the selected volume is displayed in the relationship table below the topology, as well as the lag information for all related volumes.

Snapshot

Displays the number of Snapshot copies available for a volume. Clearing the **Snapshot** check box hides all Snapshot copy information for all volumes in the topology. Clicking a Snapshot copy icon (

) displays the Snapshot copy list for a volume. The Snapshot copy count displayed next to the icon is updated approximately every hour; however, the list of Snapshot copies is updated at the

time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

Last Successful Transfer

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

History

Displays in a graph the history of incoming SnapMirror and SnapVault protection relationships for the selected volume. There are three history graphs available: incoming relationship lag duration, incoming relationship transfer duration, 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. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Protection history graphs display the following information:

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 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.



This chart is not available for volumes that are in synchronous protection relationships.

Relationship Transferred Size

Displays bytes, kilobytes, megabytes, and so on, on the vertical (y) axis depending on the transfer size, and displays days, months, or years on the horizontal (x) axis depending on the selected time period. The upper value on the y axis indicates the maximum transfer size reached in the duration period shown in the x axis. You can view the details for specific points on the graph by positioning your cursor over an area of interest.



History area

The History area displays graphs that provide information about the capacity and space reservations of the selected volume. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Graphs might be empty and the message No data found displayed when the data or the state of the volume remains unchanged for a period of time.

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

History graphs display the following information:

Volume Capacity Used

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

Volume Capacity Used vs Total

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

Volume Capacity Used (%)

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

Snapshot Capacity Used (%)

Displays the Snapshot reserve and Snapshot warning threshold as line graphs, and the capacity used by the Snapshot copies as an area graph, in percentage, on the vertical (y) axis. The Snapshot overflow is represented with different colors. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Snapshot Reserve legend, the Snapshot Reserve

graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity

Displays the severity of the event.

Event

Displays the event name.

Triggered Time

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

Related Annotations pane

The Related Annotations pane enables you to view annotation details associated with the selected volume. The details include the annotation name and the annotation values that are applied to the volume. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

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

Storage Virtual Machine

Displays the capacity and the health status of the SVM that contains the selected volume.

Aggregate

Displays the capacity and the health status of the aggregate that contains the selected volume. For FlexGroup volumes, the number of aggregates that comprise the FlexGroup is listed.

Volumes in the Aggregate

Displays the number and capacity of all the volumes that belong to the parent aggregate of the selected volume. The health status of the volumes is also displayed, based on the highest severity level. For example, if an aggregate contains ten volumes, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical. This component does not appear for FlexGroup volumes.

Qtrees

Displays the number of qtrees that the selected volume contains and the capacity of qtrees with quota that the selected volume contains. The capacity of the qtrees with quota is displayed in relation to the volume data capacity. The health status of the qtrees is also displayed, based on the highest severity level. For example, if a volume has ten qtrees, five with Warning status and the remaining five with Critical status, then the status displayed is Critical.

NFS Exports

Displays the number and status of the NFS exports associated with the volume.

CIFS Shares

Displays the number and status of the CIFS shares.

• LUNs

Displays the number and total size of all the LUNs in the selected volume. The health status of the LUNs is also displayed, based on the highest severity level.

User and Group Quotas

Displays the number and status of the user and user group quotas associated with the volume and its qtrees.

FlexClone Volumes

Displays the number and capacity of all the cloned volumes of the selected volume. The number and capacity are displayed only if the selected volume contains any cloned volumes.

Parent Volume

Displays the name and capacity of the parent volume of a selected FlexClone volume. The parent volume is displayed only if the selected volume is a FlexClone volume.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected volume.

Related Alerts pane

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

Export Policy Rules dialog box

The Export Policy Rules dialog box displays details about the export policies, authentication protocols, and access enabled on the volumes that belong to the storage virtual machine (SVM). You can use the filters to customize the display of information in the export policy rules list. By default, the information is sorted based on the index column.

Index

Displays the index assigned to the export policy rules. It is a unique number.

Access Protocols

Displays the protocols that are enabled for the export policy rules.

Client Match

Displays the clients that have permission to access data on the volumes that belong to the SVM.

Read Only Access

Displays the authentication protocol used to read data on the volumes that belong to the SVM.

Read Write Access

Displays the authentication protocol used to read or write data on the volumes that belong to the SVM.

Snapshot Copies on a Volume dialog box

You can use the Snapshot Copies on a Volume dialog box to view the list of Snapshot copies. You can delete a Snapshot copy to conserve or free disk space, or if the copy is no longer required. You can also calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted.

List view

The list view displays, in tabular format, information about the Snapshot copies on the volume. You can use the column filters to customize the data that is displayed.

Snapshot Copy

Displays the name of the Snapshot copy.

Used Space %

Displays, in percentage, the total space used by the Snapshot copy in the volume.

Total Size

Displays the total size of the Snapshot copy.

Created Time

Displays the timestamp when the Snapshot copy was created.

Dependency

Displays the applications that are dependent on the Snapshot copy. The possible values are SnapMirror, SnapVault, SnapLock, Dump, LUNs, Vclone, and Busy.

Command buttons

The command buttons enable you to perform the following tasks:

Calculate

Enables you to calculate the space that can be reclaimed by deleting one or more Snapshot copies.

Delete Selected

Deletes one or more Snapshot copies.

Close

Closes the Snapshot Copies on a Volume dialog box.

Recalculate

Enables you to calculate the space that can be reclaimed by deleting the selected Snapshot copies.

The **Recalculate** button is enabled when you make any changes in the selection of the Snapshot copies.

Managing and monitoring MetroCluster configurations

The monitoring support for MetroCluster configurations in the Unified Manager web UI enables you to check for any connectivity issues in your MetroCluster configuration. Discovering a connectivity issue early enables you to manage your MetroCluster configurations effectively.

Parts of a fabric MetroCluster configuration

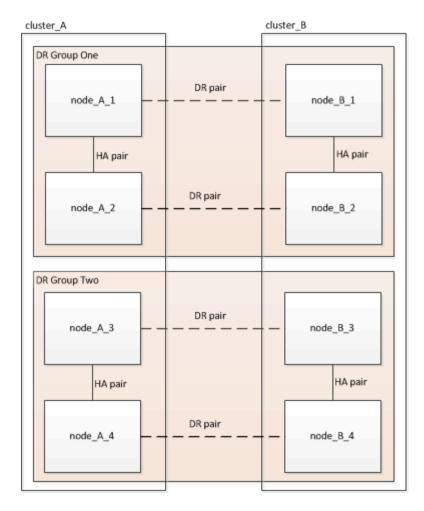
As you plan your MetroCluster configuration, you should understand the hardware components and how they interconnect.

Disaster Recovery (DR) groups

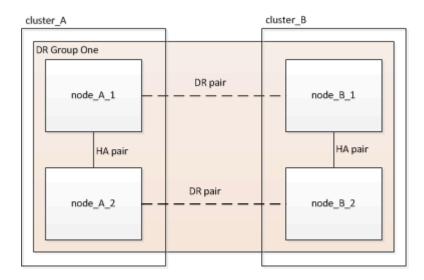
A fabric MetroCluster configuration consists of one or two DR groups, depending on the number of nodes in the MetroCluster configuration. Each DR group consists of four nodes.

- An eight-node MetroCluster configuration consists of two DR groups.
- A four-node MetroCluster configuration consists of one DR group.

The following illustration shows the organization of nodes in an eight-node MetroCluster configuration:



The following illustration shows the organization of nodes in a four-node MetroCluster configuration:



Key hardware elements

A MetroCluster configuration includes the following key hardware elements:

Storage controllers

The storage controllers are not connected directly to the storage but connect to two redundant FC switch fabrics.

• FC-to-SAS bridges

The FC-to-SAS bridges connect the SAS storage stacks to the FC switches, providing bridging between the two protocols.

· FC switches

The FC switches provide the long-haul backbone ISL between the two sites. The FC switches provide the two storage fabrics that allow data mirroring to the remote storage pools.

Cluster peering network

The cluster peering network provides connectivity for mirroring of the cluster configuration, which includes storage virtual machine (SVM) configuration. The configuration of all of the SVMs on one cluster is mirrored to the partner cluster.

Eight-node fabric MetroCluster configuration

An eight-node configuration consists of two clusters, one at each geographically separated site. cluster_A is located at the first MetroCluster site. cluster_B is located at the second MetroCluster site. Each site has one SAS storage stack. Additional storage stacks are supported, but only one is shown at each site. The HA pairs are configured as switchless clusters, without cluster interconnect switches. A switched configuration is supported, but is not shown.

An eight-node configuration includes the following connections:

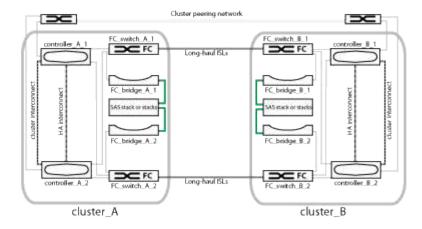
- FC connections from each controller's HBAs and FC-VI adapters to each of the FC switches
- An FC connection from each FC-to-SAS bridge to an FC switch
- SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge
- An HA interconnect between each controller in the local HA pair

If the controllers support a single-chassis HA pair, the HA interconnect is internal, occurring through the backplane, meaning that an external interconnect is not required.

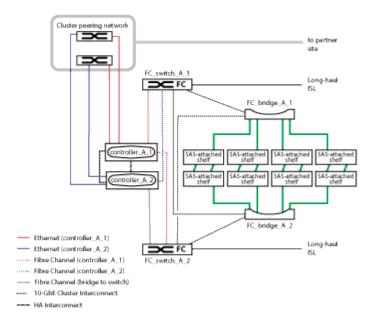
- Ethernet connections from the controllers to the customer-provided network that is used for cluster peering SVM configuration is replicated over the cluster peering network.
- · A cluster interconnect between each controller in the local cluster

Four-node fabric MetroCluster configuration

The following illustration shows a simplified view of a four-node fabric MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.

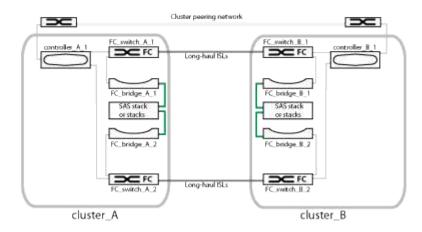


The following illustration shows a more detailed view of the connectivity in a single MetroCluster cluster (both clusters have the same configuration):



Two-node fabric MetroCluster configuration

The following illustration shows a simplified view of a two-node fabric MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.

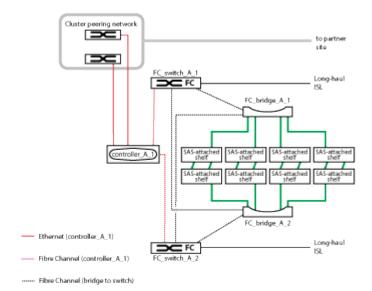


A two-node configuration consists of two clusters, one at each geographically separated site. cluster_A is located at the first MetroCluster site. cluster_B is located at the second MetroCluster site. Each site has one SAS storage stack. Additional storage stacks are supported, but only one is shown at each site.



In a two-node configuration, the nodes are not configured as an HA pair.

The following illustration shows a more detailed view of the connectivity in a single MetroCluster cluster (both clusters have the same configuration):



A two-node configuration includes the following connections:

- FC connections between the FC-VI adapter on each controller module
- FC connections from each controller module's HBAs to the FC-to-SAS bridge for each SAS shelf stack
- SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge
- Ethernet connections from the controllers to the customer-provided network that is used for cluster peering SVM configuration is replicated over the cluster peering network.

Parts of a two-node SAS-attached stretch MetroCluster configuration

The two-node MetroCluster SAS-attached configuration requires a number of parts, including two single-node clusters in which the storage controllers are directly connected to the storage using SAS cables.

The MetroCluster configuration includes the following key hardware elements:

· Storage controllers

The storage controllers connect directly to the storage using SAS cables.

Each storage controller is configured as a DR partner to a storage controller on the partner site.

· Copper SAS cables can be used for shorter distances.

Optical SAS cables can be used for longer distances.



In systems using E-Series array LUNs, the storage controllers can be directly connected to the E-Series storage arrays. For other array LUNs, connections via FC switches are required.

NetApp Interoperability Matrix Tool

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

Cluster peering network

The cluster peering network provides connectivity for mirroring of the storage virtual machine (SVM) configuration. The configuration of all SVMs on one cluster is mirrored to the partner cluster.

Parts of a two-node bridge-attached stretch MetroCluster configuration

As you plan your MetroCluster configuration, you should understand the parts of the configuration and how they work together.

The MetroCluster configuration includes the following key hardware elements:

Storage controllers

The storage controllers are not connected directly to the storage but connected to FC-to-SAS bridges. The storage controllers are connected to each other by FC cables between each controller's FC-VI adapters.

Each storage controller is configured as a DR partner to a storage controller on the partner site.

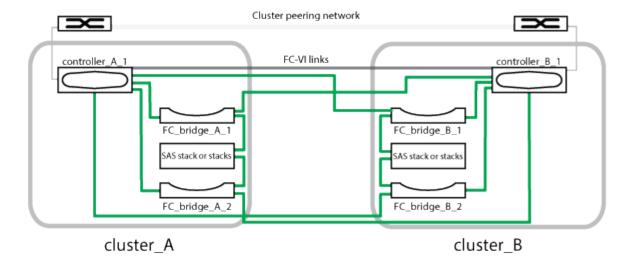
FC-to-SAS bridges

The FC-to-SAS bridges connect the SAS storage stacks to the FC initiator ports on the controllers, providing bridging between the two protocols.

· Cluster peering network

The cluster peering network provides connectivity for mirroring of the storage virtual machine (SVM) configuration. The configuration of all SVMs on one cluster is mirrored to the partner cluster.

The following illustration shows a simplified view of the MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.



- The configuration consists of two single-node clusters.
- Each site has one or more stacks of SAS storage.



SAS shelves in MetroCluster configurations are not supported with ACP cabling.

Additional storage stacks are supported, but only one is shown at each site.

Cluster connectivity status definitions

Connectivity between the clusters in a MetroCluster configuration can be one of the following statuses: Optimal, Impacted, or Down. Understanding the connectivity statuses enables you to manage your MetroCluster configurations effectively.

Connectivity status	Description	Icon displayed
Optimal	Connectivity between the clusters in the MetroCluster configuration is normal.	
Impacted	One or more errors compromise the status of failover availability; however, both of the clusters in the MetroCluster configuration are still up. For example, when the ISL link is down, when the intercluster IP link is down, or when the partner cluster is not reachable.	

Connectivity status	Description	Icon displayed
Down	Connectivity between the clusters in the MetroCluster configuration is down because one or both of the clusters are down or the clusters are in failover mode. For example, when the partner cluster is down because of a disaster or when there is a planned switchover for testing purposes.	Switchover with errors: Switchover successful:

Data mirroring status definitions

MetroCluster configurations provide data mirroring and the additional ability to initiate a failover if an entire site becomes unavailable. The status of data mirroring between the clusters in a MetroCluster configuration can either be Normal or Mirroring Unavailable. Understanding the status enables you to manage your MetroCluster configurations effectively.

Data mirroring status	Description	Icon displayed
Normal	Data mirroring between the clusters in the MetroCluster configuration is normal.	
Mirroring Unavailable	Data mirroring between the clusters in the MetroCluster configuration is unavailable because of switchover. For example, when the partner cluster is down because of a disaster or when there is a planned switchover for testing purposes.	Switchover with errors: Switchover successful:

Monitoring MetroCluster configurations

You can monitor connectivity issues in your MetroCluster configuration. The details include the status of the components and connectivity within a cluster and the connectivity status between the clusters in the MetroCluster configuration.

Before you begin

- Both the local and remote clusters in the MetroCluster configuration must be added to OnCommand Unified Manager.
- You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

You can use the information displayed in the Health/Cluster details page to rectify any connectivity issues. For example, if the connectivity between the node and the switch in a cluster is down, the following icon is displayed:



If you move the pointer over the icon, you can view detailed information about the generated event.

Unified Manager uses system health alerts to monitor the status of the components and connectivity in the MetroCluster configuration.

The MetroCluster Connectivity tab is displayed only for clusters in a MetroCluster configuration.

Steps

1. In the left navigation pane, click **Health > Clusters**.

A list of all of the monitored clusters is displayed.

- 2. From the **Health/Clusters** inventory page, click the name of the cluster for which you want to view MetroCluster configuration details.
- 3. In the Health/Cluster details page, click the MetroCluster Connectivity tab.

The topology of the MetroCluster configuration is displayed in the corresponding cluster object area.

After you finish

If you discover connectivity issues in your MetroCluster configuration, you must log in to System Manager or access the ONTAP CLI to resolve the issues.

Monitoring MetroCluster replication

You can monitor and diagnose the overall health condition of the logical connections while mirroring the data. You can identify the issues or any risk that interrupts mirroring of cluster components such as aggregates, nodes, and storage virtual machines.

Before you begin

Both the local and remote cluster in the MetroCluster configuration must be added to Unified Manager

About this task

You can use the information displayed in the Health/Cluster details page to rectify any replication issues.

If you move the pointer over the icon, you can view detailed information about the generated event.

Unified Manager uses system health alerts to monitor the status of the components and connectivity in the MetroCluster configuration.

Steps

1. In the left navigation pane, click **Health > Clusters**.

A list of the monitored clusters is displayed.

2. From the **Health/Clusters** inventory page, click the name of the cluster for which you want to view MetroCluster replication details, and then click the **MetroCluster Replication** tab.

The topology of the MetroCluster configuration to be replicated is displayed at the local site in the corresponding cluster object area with the information about the remote site where the data is being mirrored.

After you finish

If you discover mirroring issues in your MetroCluster configuration, you must log in to System Manager or access the ONTAP CLI to resolve the issues.

Managing annotations for storage objects

You can create annotations in Unified Manager to annotate storage objects. Annotations enable you to easily identify critical resources and to take appropriate actions; for example, adding critical resources to a group and assigning a group action, or creating a report of annotated resources.

What annotations are

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

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

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

How annotation rules work in Unified Manager

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

- You must associate an annotation rule to an annotation.
- You must associate an object type for an annotation rule; only one object type can be associated for an annotation rule.
- Unified Manager adds or removes annotations from storage objects after each monitoring cycle or when a rule is created, edited, deleted, or reordered.
- An annotation rule can have one or more condition groups, and each condition group can have one or more conditions.
- Storage objects can have multiple annotations. An annotation rule for a particular annotation can also use different annotations in the rule conditions to add another annotation to already annotated objects.

Conditions

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

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

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

Storage object type	Applicable operands
Volume	Object name
	Owning cluster name
	Owning SVM name
	Annotations
SVM	Object name
	Owning cluster name
	Annotations
Cluster	Object name
	Annotations

When you select annotation as an operand for any storage object, the "Is" operator is available. For all other operands, you can select either "Is" or "Contains" as operator. When you select the "Is" operator, the condition is evaluated for an exact match of the operand value with the value provided for the selected operand. When you select the "Contains" operator, the condition is evaluated to meet one of the following criteria:

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

Example of an annotation rule with conditions

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

- · Name contains "vol"
- SVM name is "data svm"

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

Condition groups

Condition groups are executed using logical OR, and then applied to storage objects. The storage objects must meet the requirements of one of the condition groups to be annotated. The storage objects that meet the conditions of all the condition groups are annotated. You can use condition groups to increase the scope of storage objects to be annotated.

Example of an annotation rule with condition groups

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

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

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

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

Description of predefined annotation values

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

· Data-priority: Mission critical

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

· Data-priority:High

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

Data-priority:Low

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

Viewing the annotation list and details

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

Steps

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

Adding annotations dynamically

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the **Annotations** page, click **Add Annotation**.
- 3. In the **Add Annotation** dialog box, type a name and description for the annotation.

You can also add values to annotations while creating annotations.

- Optional: In the Annotation Values section, click Add to add values to the annotation.
- Click Save and Close.

Adding annotations manually to individual storage objects

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

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

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

Adding values to annotations

You can add values to annotations, and then associate storage objects with a particular annotation name-value pair. Adding values to annotations helps you to manage storage objects more effectively.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot add values to predefined annotations.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the **Annotations** page, select the annotation to which you want to add a value and then click **Add** in the **Values** section.
- 3. In the Add Annotation Value dialog box, specify a value for the annotation.

The value that you specify must be unique for the selected annotation.

4. Click Add.

Creating annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Storage objects that are currently monitored are annotated as soon as the annotation rule is created. New objects are annotated only after the monitoring cycle is completed.

Steps

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

To create	Do this
A condition	a. Select an operand from the list of operands.
	b. Select either Contains or Is as the operator.
	c. Enter a value, or select a value from the available list.
A condition group	a. Click Add Condition Group.
	b. Select an operand from the list of operands.
	c. Select either Contains or Is as the operator.
	d. Enter a value, or select a value from the available list.
	e. Click Add condition to create more conditions if required, and repeat steps a through d for each condition.

7. Click Add.

Example of creating an annotation rule

Perform the following steps in the Add Annotation Rule dialog box to create an annotation rule, including configuring a condition and adding a condition group:

1. Specify a name for the annotation rule.

- 2. Select the target object type as storage virtual machine (SVM).
- 3. Select an annotation from the list of annotations, and specify a value.
- 4. In the Conditions section, select **Object Name** as the operand.
- 5. Select **Contains** as the operator.
- 6. Enter the value as svm data.
- 7. Click Add condition group.
- 8. Select **Object Name** as the operand.
- 9. Select Contains as the operator.
- 10. Enter the value as vol.
- 11. Click Add condition.
- 12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **mission-critical** as the value in step 10.
- 13. Click Add.

Configuring conditions for annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

Example of configuring a condition for an annotation rule

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

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

- 1. Enter a name for the annotation rule.
- Select the target object type as SVM.
- 3. Select an annotation from the list of annotations and a value.

- 4. In the **Conditions** field, select **Object Name** as the operand.
- 5. Select **Contains** as the operator.
- 6. Enter the value as svm data.
- 7. Click Add.

Editing annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

You cannot change the target object type for an annotation rule.

4. Click Save.

Reordering annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

1. In the toolbar, click , and then click **Annotations** in the left Management menu.

- 2. In the Annotation Rules tab, click Reorder.
- 3. In the **Reorder Annotation Rule** dialog box, drag and drop single or multiple rows to rearrange the sequence of the annotation rules.
- 4. Click Save.

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

Deleting annotations

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

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The annotation values must not be used in other annotations or group rules.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the **Annotations** tab, select the annotation that you want to delete.

The details of the selected annotation are displayed.

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

Results

The selected annotation and its value is deleted.

Deleting values from annotations

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

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The annotation value must not be associated with any annotation rules or group rules.

About this task

You cannot delete values from predefined annotations.

Steps

- 1. In the toolbar, click [, and then click **Annotations** in the left Management menu.
- 2. In the annotations list in the **Annotations** tab, select the annotation from which you want to delete a value.
- In the Values area of the Annotations tab, select the value you want to delete, and then click Delete.

4. In the Warning dialog box, click Yes.

The value is deleted and no longer displayed in the list of values for the selected annotation.

Deleting annotation rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you delete an annotation rule, the annotation is disassociated and removed from the storage objects.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the **Annotation Rules** tab, select the annotation rule that you want to delete, and then click **Delete**.
- 3. In the **Warning** dialog box, click **Yes** to confirm the deletion.

Description of Annotations windows and dialog boxes

You can view and manage all your annotations from the Management/Annotations page. You can also configure annotation rules for your storage objects from the Annotation Rules tab.

Management/Annotations page

The Management/Annotations page enables you to create annotations in Unified Manager that can be used to annotate storage objects, or you can edit or delete annotations. You can either manually annotate storage objects with an annotation=value pair or configure annotation rules. Storage objects are annotated dynamically based on the annotation you apply.

When you log in as an operator, you will have only read access to the page. You can access the add, edit, or delete buttons in each tab when you log in as a storage administrator or Unified Manager administrator.

Annotations tab

The Annotations tab enables you to view, create, edit, or delete annotations in Unified Manager.

Annotations list

Displays the names of the predefined and custom annotations. The count of the annotation values associated with each annotation is also displayed. You can click the annotation name to view the details of the annotation.

Summary area

You can view the following details of the selected annotation:

Description

Displays the description provided for the annotation.

· Created by

Displays the name of the user who created the annotation.

Creation date

Displays the date when the annotation was created.

Annotation=Values Pairs

Displays the list of annotation-value pairs and associated storage objects that are available for the selected annotation.

Value

Displays the name of the annotation=value pair.

Applicable Clusters

Displays the number of clusters that are annotated with a particular annotation=value pair. You can click the number to view the clusters page, which displays a filtered list of the clusters associated with a specific value.

Applicable storage virtual machines (SVMs)

Displays the number of storage virtual machines (SVMs) that are annotated with a particular annotation=value pair. You can click the number to view the storage virtual machines (SVMs) page, which displays a filtered list of storage virtual machines (SVMs) associated with a specific value.

Applicable Volumes

Displays the number of volumes that are annotated with a particular annotation=value pair. You can click the number to view the volumes page, which displays a filtered list of the volumes associated with a specific value.

Object Associations via Rules

Displays the list of annotation rules and the associated storage objects for the selected annotation.

• Rank

Displays the order of the annotation rules to be applied on the storage objects.

Rules

Displays the name of the annotation rule.

Target Object Type

Displays the type of storage object to which the annotation rule is applied.

Associated Annotation Value

Displays the annotation=value pair applied to the storage object.

Applicable Objects

Displays the count of the storage objects that are annotated based on the annotation rule.

Manual Object Associations

Displays the list of annotations that you have manually configured and associated with storage objects.

Annotation=Value Pair

Displays the name of the manual annotation and the value.

Applicable Clusters

Displays the number of clusters that are annotated with a particular manual annotation value. You can click the number to view the clusters page, which displays a filtered list of the clusters associated with a specific value.

Applicable storage virtual machines (SVMs)

Displays the number of storage virtual machines (SVMs) that are annotated with a particular manual annotation value. You can click the number to view the storage virtual machines (SVMs) page, which displays a filtered list of storage virtual machines (SVMs) associated with a specific value.

Applicable Volumes

Displays the number of volumes that are annotated with a particular manual annotation value. You can click the number to view the volumes page, which displays a filtered list of the volumes associated with a specific value.

Command buttons

You must have the OnCommand Administrator or Storage Administrator role. For predefined annotations, you cannot add or delete values.

Add Annotation

Opens the Add Annotation dialog box, which enables you to create new custom annotations and assign values to the annotation.

Actions

Enables you to edit or delete the selected annotation description.

• Edit

Opens the Edit Annotation dialog box, which enables you to modify the annotation name and description.

Delete

Enables you to delete the annotation value. You can delete the value only when it is not associated with any annotation rules or group rules.

Annotation Rules tab

The Annotations Rules tab displays the annotation rules you created to annotate storage objects. You can perform tasks such as adding, editing, deleting, or reordering an annotation rule. You can also view the number of storage objects that satisfy the annotation rule.

Command buttons

You must have the OnCommand Administrator or Storage Administrator role.

Add

Displays the Add Annotation Rule dialog box, which enables you to create annotation rules for storage objects.

• Edit

Displays the Edit Annotation Rule dialog box, which enables you to reconfigure previously configured annotation rules.

Delete

Deletes the selected annotation rules.

Reorder

Displays the Reorder Annotation Rule dialog box, which enables you to rearrange the order of the annotation rules.

List View

The list view displays, in tabular format, the annotation rules you created in the Unified Manager server. You can use the column filters to customize the data that is displayed. The list view of the Annotation Rules tab and the list view of the Associated Rules section in the Annotation tab contains the following columns:

- Rank
- Name
- · Target Object type
- · Associated Annotation Value
- Applicable Objects

An additional column is displayed for the Annotation Rules tab, Associated Annotation, which displays the name of the annotation applied to the storage object.

Add Annotation dialog box

The Add Annotation dialog box enables you to create custom annotations that you can

associate with clusters, volumes, and storage virtual machines (SVMs) through annotation rules.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation Name

Specifies the name of the annotation. You must enter a unique name for the annotation.

Description

Specifies a meaningful description of the annotation.

Annotation Values

Add

Adds a new value to the selected annotation.

Delete

Deletes the selected value for an annotation.

Command buttons

· Save and Close

Saves the new annotation and closes the Add Annotation dialog box dialog box.

Cancel

Closes the Add Annotation dialog box without saving your changes.

Edit Annotation dialog box

The Edit Annotation dialog box enables you to change the description of an existing annotation.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation Name

Displays the name of the annotation. This field cannot be edited.

Description

Provides a meaningful description of the annotation. You can edit this field when you want to change the current description of the annotation.

Command buttons

Save and Close

Saves the annotation description changes and closes the dialog box.

Cancel

Closes the Edit Annotation dialog box without saving your changes.

Add Annotation Rule dialog box

The Add Annotation Rule dialog box enables you to create annotation rules in Unified Manager to dynamically annotate storage objects.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Specifies the name of the annotation rule.

Target Object Type

Specifies the type of storage objects (storage virtual machines (SVMs), volumes, or clusters) that you want to annotate.

Apply Annotation

Specifies the annotation and the value you can use to annotate storage objects when all conditions are met.

Conditions

Specifies conditions that determine which storage objects you can annotate.

Command buttons

Save and Add

Adds the annotation rule you created and enables you to add another annotation rule without closing the dialog box.

Add

Adds the annotation rule and closes the Add Annotation Rule dialog box.

Cancel

Cancels the changes and closes the Add Annotation Rule dialog box.

Add Condition

Adds a condition to define the annotation rule.

Add Condition Group

Adds a condition group to define conditions for the annotation rule.

Edit Annotation Rule dialog box

You can edit the annotation rules you created to add or remove annotations on storage objects.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Displays the name of the annotation rule.

Target Object Type

Displays the type of storage object that you want to annotate. You cannot change the object type.

Apply Annotation

Displays the annotation and the value you can use to annotate storage objects when all conditions are met.

Conditions

Displays the list of conditions for the annotation rule. You can edit the conditions to add or remove the annotation on storage objects.

Command buttons

Save

Saves the changes you made and closes the Edit Annotation Rule dialog box.

Cancel

Closes the Edit Annotation Rule dialog box without saving your changes.

Reorder Annotation Rule dialog box

You can use the Reorder Annotation Rule dialog box to specify the order in which you want annotation rules to be applied to storage objects.

Command buttons

You must have the OnCommand Administrator or Storage Administrator role.

Save

Saves the changes you made to the annotation rules and closes the Reorder Annotation Rule dialog box.

Cancel

Closes the Reorder Annotation Rule dialog box without saving the changes you made.

List View

Rank

Displays the order in which the annotation rules will be applied to the storage objects.

Name

Displays the name of the annotation rule.

Target Object Type

Displays the type of storage object to which the annotation rule is applied.

Associated Annotation

Displays the name of the annotation that is applied to the storage object.

Associated Annotation Value

Displays the annotation value for the storage object.

Annotate Cluster dialog box

The Annotate Cluster dialog box enables you to manually annotate storage objects. You can select either a single cluster or multiple clusters and annotate with a specific value pair from the existing list of annotations.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation=Value Pairs

Enables you to select the required annotation for the selected cluster.

Apply

Applies the selected annotation to the cluster.

Cancel

Closes the Annotate Cluster dialog box without saving your changes.

Annotate SVM dialog box

The Annotate SVM dialog box enables you to manually annotate storage objects. You can select either a single SVM or multiple SVMs and annotate with a specific value pair from the existing list of annotations.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation=Value Pairs

Enables you to select the required annotation for the selected SVM.

Apply

Applies the selected annotation to the SVM.

Cancel

Closes the Annotate SVM dialog box without saving your changes.

Annotate Volume dialog box

The Annotate Volume dialog box enables you to manually annotate storage objects. You can select either a single volume or multiple volumes and annotate with a specific value pair from the existing list of annotations.

You must have the OnCommand Administrator or Storage Administrator role.

Annotation=Value Pairs

Enables you to select the required annotation for the selected volume.

Apply

Applies the selected annotation to the volume.

Cancel

Closes the Annotate Volume dialog box without saving your changes.

Managing and monitoring groups

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

Understanding groups

You can create groups in Unified Manager to manage storage objects. Understanding the concepts about groups and how group rules enable you to add storage objects to a group will help you to manage the storage objects in your environment.

What a group is

A group is a dynamic collection of heterogenous storage objects (clusters, SVMs, or volumes). You can create groups in Unified Manager to easily manage a set of storage objects. The members in a group might change, depending on the storage objects that are monitored by Unified Manager at a point in time.

- Each group has a unique name.
- You must configure a minimum of one group rule for each group.
- You can associate a group with more than one group rule.
- Each group can include multiple types of storage objects such as clusters, SVMs, or volumes.

- Storage objects are dynamically added to a group based on when a group rule is created or when Unified Manager completes a monitoring cycle.
- You can simultaneously apply actions on all the storage objects in a group such as setting thresholds for volumes.

How group rules work for groups

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

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

Conditions

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

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

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

Storage object type	Applicable operands
Volume	Object name
	Owning cluster name
	Owning SVM name
	Annotations
SVM	Object name
	Owning cluster name
	Annotations
Cluster	Object name
	Annotations

When you select annotation as an operand for any storage object, the "Is" operator is available. For all other operands, you can select either "Is" or "Contains" as operator.

Operand

The list of operands in Unified Manager changes based on the selected object type. The list includes the object name, owning cluster name, owning SVM name, and annotations that you define in Unified Manager.

Operator

The list of operators changes based on the selected operand for a condition. The operators supported in Unified Manager are "Is" and "Contains".

When you select the "Is" operator, the condition is evaluated for exact match of operand value to the value provided for the selected operand.

When you select the "Contains" operator, the condition is evaluated to meet one of the following criteria:

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

The value field changes based on the operand selected.

Example of a group rule with conditions

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

- · Name contains "vol"
- SVM name is "data svm"

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

Condition groups

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

Example of a group rule with condition groups

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

- Condition group 1
 - Name contains "vol"
 - SVM name is "data_svm"
 Condition group 1 selects all volumes that include "vol" in their names and that are hosted on SVMs with the name "data_svm".
- · Condition group 2
 - Name contains "vol"

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

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

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

How group actions work on storage objects

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

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

Example of unique group actions

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

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

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

Managing groups of storage objects

You can manage storage objects in your environment by creating groups of storage objects. These storage objects must satisfy the group rules associated with the group.

Adding groups

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can define group rules to add or remove members from the group and to modify group actions for the group.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- 2. In the Groups tab, click Add.
- 3. In the Add Group dialog box, enter a name and description for the group.

The group name must be unique.

4. Click Add**.

Deleting groups

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

Before you begin

- None of the storage objects (clusters, SVMs, or volumes) must be associated with any group rule that is associated with the group that you want to delete.
- You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

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

Editing groups

You can edit the name and description of a group that you created in Unified Manager.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you edit a group to update the name, you must specify a unique name; you cannot use an existing group name.

Steps

- In the toolbar, click , and then click Management > Groups.
- In the Groups tab, select the group that you want to edit, and then click Edit.

- 3. In the **Edit Group** dialog box, change the name, description, or both for the group.
- 4. Click Save.

Adding group rules

You can create group rules for a group to dynamically add storage objects such as volumes, clusters, or storage virtual machines (SVMs) to the group. You must configure at least one condition group with at least one condition to create a group rule.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Storage objects that are currently monitored are added as soon as the group rule is created. New objects are added only after the monitoring cycle is completed.

Steps

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

To create	Do this
A condition	a. Select an operand from the list of operands.
	b. Select either Contains or Is as the operator.
	c. Enter a value, or select a value from the available list.
A condition group	a. Click Add Condition Group
	b. Select an operand from the list of operands.
	c. Select either Contains or Is as the operator.
	d. Enter a value, or select a value from the available list.
	e. Click Add condition to create more conditions if required, and repeat steps a through d for each condition.

7. Click Add.

Example for creating a group rule

Perform the following steps in the Add Group Rule dialog box to create a group rule, including configuring a condition and adding a condition group:

- 1. Specify a name for the group rule.
- Select the object type as storage virtual machine (SVM).
- 3. Select a group from the list of groups.
- 4. In the Conditions section, select **Object Name** as the operand.
- 5. Select **Contains** as the operator.
- 6. Enter the value as svm_data.
- Click Add condition group.
- 8. Select **Object Name** as the operand.
- 9. Select **Contains** as the operator.
- 10. Enter the value as vol.
- 11. Click Add condition.
- 12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **critical** as the value in step 10.
- 13. Click **Add** to create the condition for the group rule.

Editing group rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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



You cannot change the target object type for a group rule.

4. Click Save.

Deleting group rules

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When a group rule is deleted, the associated storage objects will be removed from the group.

Steps

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

Configuring conditions for group rules

You can configure one or more conditions to create group rules in Unified Manager that are applied on the storage objects. The storage objects that satisfy the group rule are combined into a group.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- 2. Click Add.
- In the Add Group Rule dialog box, select one operand from the list of operands.
- 4. Select an operator for the condition.
- 5. Enter a required value or select one from the available list.
- 6. Click Add.

Example of configuring a condition for a group rule

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

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

- 1. Enter a name for the group rule.
- Select the object type as SVM.
- 3. Select a group from the list of groups.
- 4. In the **Conditions** field, select **Object Name** as the operand.
- 5. Select **Contains** as the operator.
- 6. Enter the value as svm data.
- 7. Click Add.

Adding group actions

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- In the Group Actions tab, click Add.
- 3. In the Add Group Action dialog box, enter a name and description for the action.
- 4. From the **Group** menu, select a group for which you want to configure the action.
- 5. From the **Action Type** menu, select an action type.

The dialog box expands, enabling you to configure the selected action type with required parameters.

- 6. Enter appropriate values for the required parameters to configure a group action.
- 7. Click Add.

Editing group actions

You can edit the group action parameters that you configured in Unified Manager, such as the group action name, description, associated group name, and parameters of the action type.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

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

Configuring volume health thresholds for groups

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

Steps

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

Deleting group actions

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

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you delete the group action for the volume health threshold, global thresholds are applied to the storage objects in that group. Any object-level health thresholds that are set on the storage object are not impacted.

Steps

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

Reordering group actions

You can change the order of the group actions that are to be applied to the storage objects in a group. Group actions are applied to storage objects sequentially based on their rank. The lowest rank is assigned to the group action that you configured last. You can change the rank of the group action depending on your requirements.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can select either a single row or multiple rows, and then perform multiple drag-and-drop operations to change the rank of group actions. However, you must save the changes for the re-prioritization to be reflected

in the group actions grid.

Steps

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

Description of groups windows and dialog boxes

You can use the Management/Groups page to view and manage all your groups. You can also configure group rules and actions for your storage objects from the Group Rules tab and Group Actions tab.

Management/Groups page

The Management/Groups page enables you to create groups in Unified Manager to easily manage storage objects. A group is a dynamic collection of storage objects (clusters, volumes, and SVMs), which is defined by the group rules you create for the group.

The Management/Groups page includes tabs that enable you to add, delete, or edit a group, group rules, and group actions. When you log in as an operator, you will have only read access to the page. You can access the add, edit, or delete buttons in each tab when you log in as a storage administrator or Unified Manager administrator.

Groups tab

The Groups tab displays the name and description of the groups you created. You can perform tasks such as adding, editing, or deleting a group. The tab also displays the number of group rules and group actions associated with a group, the number of clusters, SVMs, and volumes in the group.

Command buttons

Add

Displays the Add Group dialog box, which enables you to add a group and provide a name and description for the group.

You can also apply group rules later to the group to include storage objects.

• Edit

Displays the Edit Group dialog box, which enables you to edit the name and description for the selected group.

Delete

Deletes the selected group.

List view

The list view displays, in tabular format, the groups you created in Unified Manager. You can use the column filters to customize the data that is displayed. By default, the list is sorted by group name.

Name

Displays the name of the group.

Description

Displays the description of the group.

Associated Rules

Displays the number of rules added to the group.

Associated Actions

Displays the number of group actions added to the group.

Applicable Clusters

Displays the number of clusters included in the group.

Applicable SVMs

Displays the number of SVMs included in the group.

Applicable Volumes

Displays the number of volumes included in the group.

Group Rules tab

The Group Rules tab displays the group rules you created for groups to contain storage objects. You can perform tasks such as adding, editing, or deleting a group rule. The tab also displays the group name for which the group rule is created and the storage object for which the rule is applied. You can also view the number of storage objects that satisfy the group rule.

Command buttons

Add

Displays the Add Group Rule dialog box, which enables you to create group rules for storage objects.

• Edit

Displays the Edit Group Rule dialog box, which enables you to reconfigure previously configured group rules.

Delete

Deletes the selected group rule.

List view

The list view displays, in tabular format, the group rules you created for a specific storage object (clusters, volumes, or SVMs) and the count of storage objects that satisfy the defined group rule.

Name

Displays the name of the rule.

Associated Group

Displays the name of the group for which the group rule is defined.

Target Object Type

Displays the type of storage object to which the group rule is applied.

Applicable Objects

Displays the count of the storage objects included in the group based on the group rule.

Group Actions tab

The Group Actions tab displays the name and type of group actions you define for groups. You can perform tasks such as adding, editing, deleting, or reordering the group actions. The tab also displays the name of the group on which the group action is applied.

Command buttons

Add

Displays the Add Action dialog box, which enables you to create group actions for a group of storage objects. For example, you can set the threshold levels of storage objects in a group.

• Edit

Displays the Edit Action dialog box, which enables you to reconfigure previously configured group actions.

Delete

Deletes the selected group action.

Reorder

Displays the Reorder Group Actions dialog box to rearrange the order of the group actions.

List view

The list view displays, in tabular format, the group actions you created for the groups in the Unified Manager server. You can use the column filters to customize the data that is displayed.

Rank

Displays the order of the group actions to be applied on the storage objects in a group.

Name

Displays the name of the group action.

Associated Group

Displays the name of the group for which the group action is defined.

Action Type

Displays the type of group action that you can perform on the storage objects in a group.

You cannot create multiple group actions of the same action type for a group. For example, you can create a group action of setting volume thresholds for a group. However, you cannot create another group action for the same group to change volume thresholds.

Description

Displays the description of the group action.

Add Group dialog box

The Add Group dialog box enables you to create groups to include clusters, volumes, and SVMs based on the group rules.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Specifies the name of the group. You must enter a unique name for the group.

Description

Specifies a meaningful description of the group.

Command buttons

The command buttons enable you to add or cancel the creation of a new group.

Add

Creates the new group.

Cancel

Closes the Add Group dialog box without saving your changes.

Edit Group dialog box

The Edit Group dialog box enables you to change the name and description of a group.

You must have the OnCommand Administrator or Storage Administrator role.

Group Name

Displays the name of the group. When changing the group name, you must not use an existing group name.

Description

Provides a meaningful description of the group. You can edit this field when you want to change the current description of the group.

Command buttons

The command buttons enable you to save or cancel changes you make to the group.

Save

Saves the changes you made and closes the dialog box.

Cancel

Closes the Edit Group dialog box without saving your changes.

Groups details page

From the Groups details page, you can view the details of a selected group. You can also view additional information such as the group rules and group actions associated with the selected group.

Command buttons

View Groups

Enables you to navigate to the Groups page.

Actions

Enables you to edit or delete the group, based on your role. You must have the OnCommand Administrator or Storage Administrator role.

Manage Group Rules

Enables you to navigate to the Group Rules page, which displays rules for this group.

Manage Group Actions

Enables you to navigate to the Group Actions page, which displays actions for this group.

Summary area

You can view the following group details:

Description

Displays the description provided for the group.

· Created by

Displays the name of the user who created the group.

Creation Date

Displays the date when the group was created.

Associated Rules

Displays all the group rules created for a group, in tabular format. You can view the details of each group rule, such as the rule name, associated object type, and the count of storage objects of the associated object type.

Associated Actions

Displays all the group actions, configured for a group, in tabular format. You can view the details of each group action, such as the rank, name, action type, and description.

Add Group Rule dialog box

The Add Group Rule dialog box enables you to create group rules in Unified Manager to dynamically group storage objects. You can later configure and apply group actions for the group.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Specifies the name of the group rule.

Target Object Type

Specifies the type of storage objects to include in the group.

Group

Specifies the name of the group for which the group rule is created.

Conditions

Specifies conditions that determine which storage objects can be included in a group.

Condition group

Specifies condition groups which have one or more conditions defined for including storage objects in a group.

Command buttons

Save and Add

Adds the group rule and enables you to add another group rule without closing the dialog box.

Add

Adds the group rule and closes the Add Group Rule dialog box.

Cancel

Cancels the changes and closes the Add Group Rule dialog box.

Add Condition

Adds a condition to define the group rule.

Add Condition Group

Adds a condition group to define conditions for the group rule.

Edit Group Rule dialog box

You can edit the group rules you created to include the maximum number of storage objects in a group.

You must have the OnCommand Administrator or Storage Administrator role.

Rule Name

Displays the name of the rule.

Target Object Type

Displays the storage object to be added to a selected group. You cannot change the object type.

Associated Group

Displays the associated group. You can select a different group for the group rule.

Condition

Displays the list of conditions for a selected group. You can edit the conditions. The storage objects are either removed or added to a selected group based on the changes.

Command buttons

Save

Saves the changes you made and closes the dialog box.

Cancel

Closes the Edit Group Rule dialog box without saving your changes.

Add Group Action dialog box

The Add Group Action dialog box enables you to configure group actions that can be

applied to storage objects of a selected group.

You must have the OnCommand Administrator or Storage Administrator role.

Name

Specifies the name of the action.

Description

Specifies the description of the action.

Group

Specifies the group for which the action is configured.

Action type

Specifies the type of action configured. Based on the selected action type, the Add Group Action dialog box expands, enabling you to configure a group action by providing the required values.

Unified Manager currently supports only volume threshold action type.

Command buttons

Add

Adds the new action and closes the dialog box.

Cancel

Closes the Add Group Action dialog box dialog box without saving your changes.

Group action-volume thresholds section

The group action-volume thresholds section enables you to configure group-level health thresholds for volumes. These thresholds are applied to all the volumes in a group. When the volume health thresholds are configured at the group level, the global health threshold values are not affected.

You can configure volume health thresholds for the following to configure a group action:

- · Capacity
- Growth
- · Qtree quota
- · Snapshot copies
- Inodes

Global default values are used if volume health thresholds are not configured for any of these categories. You can set health thresholds for the following:

Capacity

- Growth
- · Qtree quota
- · Snapshot copies
- Inodes

Capacity section

You can set conditions for the following volume capacity health thresholds:

Space Nearly Full

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

· Default value: 80 percent

The value for this threshold must be lower than the value for the Volume Full threshold for the management server to generate an event.

· Event generated: Volume Nearly Full

Event severity: Warning

Space Full

Specifies the percentage at which a volume is considered full:

Default value: 90 percent

Event generated: Volume Full

· Event severity: Error

Overcommitted

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

Default value: 100 percent

Event generated: Volume Overcommitted

· Event severity: Error

Growth section

You can set the following health threshold conditions for volume growth:

Growth Rate

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

Default value: 1 percent

Event generated: Volume Growth Rate Abnormal

· Event severity: Warning

Growth Rate Sensitivity

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

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

o Default value: 2

Qtree Quota section

You can set the following health threshold conditions for volume quotas:

Nearly Overcommitted

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

Default value: 95 percent

Event generated: Volume Qtree Quota Nearly Overcommitted

Event severity: Warning

Overcommitted

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

· Default value: 100 percent

Event generated: Volume Qtree Quota Overcommitted

· Event severity: Error

Snapshot Copies section

You can set the following health threshold conditions for the Snapshot copies in the volume:

Snapshot Reserve Full

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

Default value: 90 percent

Event generated: Volume Snapshot Reserve Full

Event severity: Error

· Days Until Full

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

Default value: 7

Event generated: Volume Snapshot Reserve Days Until Full

Event severity: Error

Count

Specifies the number of Snapshot copies on a volume that are considered to be too many:

Default value: 250

Event generated: Too Many Snapshot Copies

· Event severity: Error

Inodes section

You can set the following health threshold conditions for inodes:

Nearly Full

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

· Default value: 80 percent

Event generated: Inodes Nearly Full

Event severity: Warning

• Full

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

· Default value: 90 percent

Event generated: Inodes Full

· Event severity: Error

Edit Group Action dialog box

You can edit the group action that you created for groups by using the Edit Group Action dialog box.

You must have the OnCommand Administrator or Storage Administrator role.

Action Name

Displays the name of the group action.

Description

Displays the description of the group action.

Group

Displays the name of the group selected.

Action type

Displays the type of group action. You cannot change the action type. However, you can modify the parameters that you used to configure the group action.

Command buttons

Save

Saves the changes you made to the group action.

Cancel

Closes the Edit Group Action dialog box without saving your changes.

Reorder Group Actions dialog box

You can use the Reorder Group Actions dialog box to change the ranks of one or more group actions. The position of a group action in the grid determines the rank for the group action.

You must have the OnCommand Administrator or Storage Administrator role.

Rank

Specifies the order of the group action to be applied on storage objects in a group.

Name

Specifies the name of the group action.

Action Type

Specifies the type of action that you can perform on the storage objects in a group.

Associated Group

Specifies the name of the group for which the group actions are defined.

Managing and monitoring protection relationships

OnCommand Unified Manager enables you to create protection relationships, to monitor and troubleshoot SnapMirror and SnapVault relationships on managed clusters, and to restore data when it is overwritten or lost.

For SnapMirror operations there are two replication types:

Asynchronous

Replication from the primary to the secondary volume is determined by a schedule.

• Synchronous

Replication is performed simultaneously on the primary and secondary volume.

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

What resource pools are

Resource pools are groups of aggregates that are created by a storage administrator using Unified Manager to provide provisioning to partner applications for backup management.

You might pool your resources based on attributes such as performance, cost, physical location, or availability. By grouping related resources into a pool, you can treat the pool as a single unit for monitoring and provisioning. This simplifies the management of these resources and allows for a more flexible and efficient use of the storage.

During secondary storage provisioning, Unified Manager determines the most suitable aggregate in the resource pool for protection using the following criteria:

- The aggregate is a data aggregate (not a root aggregate) and it is ONLINE.
- The aggregate is on a destination cluster node whose ONTAP version is the same or greater than the source cluster major version.
- The aggregate has the largest available space of all the aggregates in the resource pool.
- After provisioning the destination volume, the aggregate space is within the nearly-full and nearly overcommitted threshold defined for the aggregate (global or local threshold, whichever is applicable).
- The number of FlexVol volumes on the destination node must not exceed the platform limit.

Types of SnapMirror protection

Depending on the deployment of your data storage topology, Unified Manager enables you to configure multiple types of SnapMirror protection relationships. All variations of SnapMirror protection offer failover disaster recovery protection, but offer differing capabilities in performance, version flexibility, and multiple backup copy protection.

Traditional SnapMirror Asynchronous protection relationships

Traditional SnapMirror Asynchronous protection provides block replication mirror protection between source and destination volumes.

In traditional SnapMirror relationships, mirror operations execute faster than they would in alternative SnapMirror relationships because the mirror operation is based on block replication. However, traditional SnapMirror protection requires that the destination volume run under the same or later minor version of ONTAP software as the source volume within the same major release (for example, version 8.x to 8.x, or 9.x to 9.x).

SnapMirror Asynchronous protection with version-flexible replication

SnapMirror Asynchronous protection with version-flexible replication provides logical replication mirror protection between source and destination volumes, even if those volumes are running under different versions of ONTAP 8.3 or later software (for example, version 8.3 to 8.3, or 8.3 to 9.1, or 9.0 to 8.3).

In SnapMirror relationships with version-flexible replication, mirror operations do not execute as quickly as they would in traditional SnapMirror relationships.

Because of slower execution, SnapMirror with version-flexible replication protection is not suitable to implement in either of the following circumstances:

- The source object contains more than 10 million files to protect.
- The recovery point objective for the protected data is two hours or less. (That is, the destination must always contain mirrored, recoverable data that is no more than two hours older than data at the source.)

In either of the listed circumstances, the faster block-replication based execution of default SnapMirror protection is required.

SnapMirror Asynchronous protection with version-flexible replication and backup option

SnapMirror Asynchronous protection with version-flexible replication and backup option provides mirror protection between source and destination volumes and the capability to store multiple copies of the mirrored data at the destination.

The storage administrator can specify which Snapshot copies are mirrored from source to destination and can also specify how long to retain those copies at the destination, even if they are deleted at the source.

In SnapMirror relationships with version-flexible replication and backup option, mirror operations do not execute as guickly as they would in traditional SnapMirror relationships.

SnapMirror Synchronous protection with strict synchronization

SnapMirror Synchronous protection with "strict" synchronization ensures that the primary and secondary volumes are always a true copy of each other. If a replication failure occurs when attempting to write data to the secondary volume, then the client I/O to the primary volume is disrupted.

SnapMirror Synchronous protection with regular synchronization

SnapMirror Synchronous protection with "regular" synchronization does not require that the primary and secondary volume are always a true copy of each other; thereby ensuring availability of the primary volume. If a replication failure occurs when attempting to write data to the secondary volume, the primary and secondary volumes fall out of sync and client I/O will continue to the primary volume.



The Restore button and the Relationship operation buttons are not available when monitoring synchronous protection relationships from the Health/Volumes inventory page or the Health/Volume details page.

Viewing volume protection relationships

From the Protection/Volume Relationships page, you can view the status of existing volume SnapMirror and SnapVault relationships. You can also examine details about protection relationships, including transfer and lag status, source and destination details, schedule and policy information, and so on.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can also initiate relationship commands from this page.

Steps

1. In the left navigation pane, click **Protection > Volume Relationships**.

The Protection/Volume Relationships page is displayed.

- 2. Choose one of the following ways to view the volume protection details:
 - To view current information about all the volume relationships, remain on the default All Volume Relationships page.
 - To view detailed information about the volume transfer trends over a period of time, in the View menu, select **Volume Transfer Status (Historical)**.
 - To view detailed information about the volume transfer activity on a day to day basis, in the View menu, select **Volume Transfer Rate (Historical)**.



The volume transfer views display information for volumes in asynchronous relationships only - volumes in synchronous relationships are not shown.

Creating a SnapVault protection relationship from the Health/Volumes inventory page

You can use the Health/Volumes inventory page to create SnapVault relationships for one or more volumes on the same storage virtual machine (SVM) to enable data backups for protection purposes.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

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

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

Steps

- 1. In the left navigation pane, click **Health > Volumes**.
- 2. In the **Health/Volumes** inventory page, right-click a volume you want to protect and select **Protect**.

Alternatively, to create multiple protection relationships on the same storage virtual machine (SVM), select one or more volumes in the Health/Volumes inventory page, and click **Protect** on the toolbar.

3. Select **SnapVault** from the menu.

The Configure Protection dialog box is launched.

- 4. Click SnapVault to view the SnapVault tab and to configure the secondary volume information.
- 5. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then

click Apply.

- 6. Complete the **Destination Information** area and the **Relationship Settings** area in the **SnapVault** tab.
- Click Apply.

You are returned to the Health/Volumes inventory page.

8. Click the protection configuration job link at the top of the **Health/Volumes** inventory page.

If you are creating only one protection relationship, the Protection/Job details page is displayed; however, if you are creating more than one protection relationship, a filtered list of all the jobs associated with the protection operation is displayed.

- 9. Do one of the following:
 - If you have only one job, click Refresh to update the tasks list and task details associated with the
 protection configuration job and to determine when the job is complete.
 - If you have more than one job:
 - i. Click a job in the jobs list.
 - ii. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
 - iii. Use the **Back** button to return to the filtered list and view another job.

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

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

Before you begin

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

About this task

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

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

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, right-click a volume in the topology view that you want to protect.
- 2. Select **Protect > SnapVault** from the menu.

The Configure Protection dialog box is launched.

- 3. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary resource information.
- 4. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then

click Apply.

- Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
- 6. Click Apply.

You are returned to the Health/Volume details page.

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

The Protection/Job details page is displayed.

8. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

When the job tasks are complete, the new relationships are displayed in the Health/Volume details page topology view.

Creating a SnapMirror protection relationship from the Health/Volumes inventory page

Using the Health/Volumes inventory page enables you to create several SnapMirror protection relationships at one time by selecting more than one volume on the same SVM.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

About this task

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

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

Steps

1. In the Health/Volumes inventory page, right-click a volume that you want to protect.

Alternatively, to create multiple protection relationships on the same SVM, select one or more volumes in the Health/Volumes inventory page, and click **Protect** > **SnapMirror** on the toolbar.

The Configure Protection dialog box is displayed.

- 2. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
- Click Advanced to set the space guarantee, as needed, and then click Apply.
- 4. Complete the **Destination Information** area and the **Relationship Settings** area in the **SnapMirror** tab.
- 5. Click Apply.

You are returned to the Health/Volumes inventory page.

6. Click the protection configuration job link at the top of the **Health/Volumes** inventory page.

If you are creating only one protection relationship, the Protection/Job details page is displayed; however, if you are creating more than one protection relationship, a list of all the jobs associated with the protection operation is displayed.

7. Do one of the following:

- If you have only one job, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
- If you have more than one job:
 - i. Click a job in the jobs list.
 - ii. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
 - iii. Use the **Back** button to return to the filtered list and view another job.

Results

Depending on the destination SVM you specified during configuration or on the options you enabled in your Advanced settings, the resulting SnapMirror relationship might be one of several possible variations:

- If you specified a destination SVM that runs under the same or a newer version of ONTAP compared to that of the source volume, a block-replication-based SnapMirror relationship is the default result.
- If you specified a destination SVM that runs under the same or a newer version of ONTAP (8.3 or later) compared to that of the source volume, but you enabled version-flexible replication in the Advanced settings, a SnapMirror relationship with version-flexible replication is the result.
- If you specified a destination SVM that runs under an earlier version of ONTAP 8.3 or a later version than that of the source volume, and the earlier version supports version-flexible replication, a SnapMirror relationship with version-flexible replication is the automatic result.

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

You can use the Health/Volume details page to create a SnapMirror relationship so that data replication is enabled for protection purposes. SnapMirror replication enables you to restore data from the destination volume in the event of data loss on the source.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

About this task

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

- · If RBAC settings do not allow this action: for example, if you have only operator privileges
- If the volume is a FlexGroup volume
- When the volume ID is unknown: for example, when you have an intercluster relationship and the

destination cluster has not yet been discovered

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

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, right-click in the topology view the name of a volume that you want to protect.
- 2. Select **Protect > SnapMirror** from the menu.

The Configure Protection dialog box is displayed.

- 3. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
- 4. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
- Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
- 6. Click Apply.

You are returned to the Health/Volume details page.

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

The job's tasks and details are displayed in the Protection/Job details page.

- 8. In the **Protection/Job** details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
- 9. When the job tasks are complete, click **Back** on your browser to return to the **Health/Volume** details page.

The new relationship is displayed in the Health/Volume details page topology view.

Results

Depending on the destination SVM you specified during configuration or on the options you enabled in your Advanced settings, the resulting SnapMirror relationship might be one of several possible variations:

- If you specified a destination SVM that runs under the same or a newer version of ONTAP compared to that of the source volume, a block-replication-based SnapMirror relationship is the default result.
- If you specified a destination SVM that runs under the same or a newer version of ONTAP (version 8.3 or higher) compared to that of the source volume, but you enabled version-flexible replication in the Advanced settings, a SnapMirror relationship with version-flexible replication is the result.
- If you specified a destination SVM that runs under an earlier version of ONTAP 8.3, or a version that is higher than that of the source volume and the earlier version supports version-flexible replication, a SnapMirror relationship with version-flexible replication is the automatic result.

Creating a SnapMirror relationship with version-flexible replication

You can create a SnapMirror relationship with version-flexible replication. Version-flexible replication enables you to implement SnapMirror protection even if source and

destination volumes run under different versions of ONTAP.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.
- The source and destination SVMs must each have a SnapMirror license enabled.
- The source and destination SVMs must each run under a version of ONTAP software that supports version-flexible replication.

About this task

SnapMirror with version-flexible replication enables you to implement SnapMirror protection even in heterogeneous storage environments in which not all storage is running under one version of ONTAP; however, mirror operations performed under SnapMirror with version-flexible replication do not execute as quickly as they would under traditional block replication SnapMirror.

Steps

- 1. Display the **Configure Protection** dialog box for the volume that you want to protect.
 - If you are viewing the Protection tab of the Health/Volume details page, right-click in the topology view that has the name of a volume that you want to protect and select **Protect > SnapMirror** from the menu.
 - If you are viewing the Health/Volumes inventory page, locate a volume that you want to protect and right-click it; then select **Protect** > **SnapMirror** from the menu.
 The Configure Protection dialog box is displayed.
- 2. Click **SnapMirror** to view the **SnapMirror** tab.
- 3. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.

If you specify a destination SVM that runs under an earlier version of ONTAP than the source volume you are protecting, and if that earlier version supports version-flexible replication, this task automatically configures SnapMirror with version-flexible replication.

- 4. If you specify a destination SVM that runs under the same version of ONTAP as that of the source volume, but you still want to configure SnapMirror with version-flexible replication, click **Advanced** to enable version-flexible replication and then click **Apply**.
- 5. Click Apply.

You are returned to the Health/Volume details page.

6. Click the protection configuration job link at the top of the **Health/Volume** details page.

The jobs tasks and details are displayed in the Protection/Job details page.

- 7. In the **Protection/Job** details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
- 8. When the job tasks are complete, click **Back** on your browser to return to the **Health/Volume** details page.

The new relationship is displayed in the Health/Volume details page topology view.

Creating SnapMirror relationships with version-flexible replication with backup option

You can create a SnapMirror relationship with version-flexible replication and backup option capability. Backup option capability enables you to implement SnapMirror protection and also retain multiple versions of backup copies at the destination location.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.
- The source and destination SVMs must each have a SnapMirror license enabled.
- The source and destination SVMs must each have a SnapVault license enabled.
- The source and destination SVMs must each run under a version of ONTAP software (8.3 or higher) that supports version-flexible replication.

About this task

Configuring SnapMirror with backup option capability enables you to protect your data with SnapMirror disaster recovery capabilities, such as volume failover ability, and at the same time provide SnapVault capabilities, such as multiple backup copy protection.

Steps

- 1. Display the **Configure Protection** dialog box for the volume that you want to protect.
 - If you are viewing the Protection tab of the Health/Volume details page, right-click in the topology view the name of a volume that you want to protect and select **Protect** > **SnapMirror** from the menu.
 - If you are viewing the Health/Volumes inventory page, locate a volume you want to protect and rightclick it; then select **Protect** > **SnapMirror** from the menu.
 The Configure Protection dialog box is displayed.
- 2. Click **SnapMirror** to view the **SnapMirror** tab.
- Complete the Destination Information area and the Relationship Settings area in the Configure Protection dialog box.
- 4. Click Advanced to display the Advanced Destination Settings dialog box.
- 5. If the **Version-Flexible Replication** check box is not already selected, select it now.
- 6. Select the With backup option check box to enable backup option capability; then click Apply.
- 7. Click Apply.

You are returned to the Health/Volume details page.

8. Click the protection configuration job link at the top of the **Health/Volume** details page.

The jobs tasks and details are displayed in the Protection/Job details page.

- 9. In the **Protection/Job** details page, click **Refresh** to update the task list and task details associated with the protection configuration job and to determine when the job is complete.
- 10. When the job tasks are complete, click **Back** on your browser to return to the **Health/Volume** details page.

The new relationship is displayed in the Health/Volume details page topology view.

Configuring destination efficiency settings

You can configure destination efficiency settings such as deduplication, compression, autogrow, and space guarantee on a protection destination using the Advanced Destination Settings dialog box. You use these settings when you want to maximize space utilization on a destination or secondary volume.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

By default, efficiency settings match those of the source volume, except for compression settings in a SnapVault relationship, which are disabled by default.

Steps

- 1. Click either the **SnapMirror** tab or the **SnapVault** tab in the **Configure Protection** dialog box, depending on the type of relationship you are configuring.
- 2. Click Advanced in the Destination Information area.

The Advanced Destination Settings dialog box is opened.

- 3. Enable or disable the efficiency settings for deduplication, compression, autogrow, and space guarantee, as required.
- 4. Click Apply to save your selections and return to the Configure Protection dialog box.

Creating SnapMirror and SnapVault schedules

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

Before you begin

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

Steps

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

The Create Schedule dialog box is displayed.

- 2. In the **Schedule Name** field, type the name you want to give to the schedule.
- 3. Select one of the following:
 - Basic

Select if you want to create a basic interval-style schedule.

Advanced

Select if you want to create a cron-style schedule.

4. Click Create.

The new schedule is displayed in the SnapMirror Schedule or SnapVault Schedule drop-down list.

Creating cascade or fanout relationships to extend protection from an existing protection relationship

You can extend protection from an existing relationship by creating either a fanout from the source volume or a cascade from the destination volume of an existing relationship. You might do this when you need to copy data from one site to many sites or to provide additional protection by creating more backups.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. Click **Protection > Volume Relationships**.
- 2. From the **Protection/Volume Relationships** page, select the SnapMirror relationship from which you want to extend protection.
- 3. On the action bar, click **Extend Protection**.
- 4. In the menu, select either **From Source** or **From Destination**, depending on whether you are creating a fanout relationship from the source or a cascade relationship from the destination.
- 5. Select either **With SnapMirror** or **With SnapVault**, depending on the type of protection relationship you are creating.

The Configure Protection dialog box is displayed.

6. Complete the information as indicated in the **Configure Protection** dialog box.

Editing protection relationships from the Protection/Volume Relationships page

You can edit existing protection relationships to change the maximum transfer rate, the protection policy, or the protection schedule. You might edit a relationship to decrease the bandwidth used for transfers, or to increase the frequency of scheduled transfers because data is changing often.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The selected volumes must be protection relationship destinations. You cannot edit relationships when source volumes, load-sharing volumes, or volumes that are not the destination of a SnapMirror or SnapVault relationship are selected.

Steps

1. From the **Protection/Volume Relationships** page, select in the volumes list one or more volumes in the same SVM for which you want to edit relationship settings, and then select **Edit** from the toolbar.

The Edit Relationship dialog box is displayed.

- 2. In the **Edit Relationship** dialog box, edit the maximum transfer rate, protection policy, or protection schedule, as needed.
- 3. Click Apply.

The changes are applied to the selected relationships.

Editing protection relationships from the Health/Volume details page

You can edit existing protection relationships to change the current maximum transfer rate, protection policy, or protection schedule. You might edit a relationship to decrease the bandwidth used for transfers, or to increase the frequency of scheduled transfers because data is changing often.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have installed and configured Workflow Automation.

About this task

The selected volumes must be protection relationship destinations. You cannot edit relationships when source volumes, load-sharing volumes, or volumes that are not the destination of a SnapMirror or SnapVault relationship are selected.

Steps

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

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

The Edit Relationship dialog box is displayed.

- 3. In the **Edit Relationship** dialog box, edit the maximum transfer rate, protection policy, or protection schedule, as needed.
- Click Apply.

The changes are applied to the selected relationships.

Creating a SnapMirror policy to maximize transfer efficiency

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

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.
- This task assumes that you have already completed the Destination Information area in the Configure Protection dialog box.

Steps

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

The Create SnapMirror Policy dialog box is displayed.

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

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

Creating a SnapVault policy to maximize transfer efficiency

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

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.
- You must have already completed Destination Information area in the Configure Protection dialog box.

Steps

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

The SnapVault tab is displayed.

- 2. In the **Policy Name** field, type the name that you want to give the policy.
- In the Transfer Priority field, select the transfer priority that you want to assign to the policy.
- 4. 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.

Aborting an active data protection transfer from the Protection/Volume Relationships page

You can abort an active data protection transfer when you want to stop a SnapMirror replication that is in progress. You can also clear the restart checkpoint for transfers subsequent to the baseline transfer. You might abort a transfer when it conflicts with another operation, such as a volume move.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role...
- You must have set up Workflow Automation.

About this task

The abort action does not display in the following instances:

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

You cannot clear the restart checkpoint for a baseline transfer.

Steps

1. To abort transfers for one or more protection relationships, from the **Protection/Volume Relationships** page, select one or more volumes and, on the toolbar, click **Abort**.

The Abort Transfer dialog box is displayed.

- If you want to clear the restart checkpoint for a transfer that is not a baseline transfer, select Clear Checkpoints.
- 3. Click Continue.

The Abort Transfer dialog box is closed, and the status of the abort job displays at the top of the

Protection/Volume Relationships page, along with a link to the job details.

4. Click the **View details** link to go to the **Protection/Job** details page for additional details and to view job progress.

Aborting an active data protection transfer from the Health/Volume details page

You can abort an active data protection transfer when you want to stop a SnapMirror replication that is in progress. You can also clear the restart checkpoint for a transfer if it is not a baseline transfer. You might abort a transfer when it conflicts with another operation, such as a volume move.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role..
- · You must have set up Workflow Automation.

About this task

The abort action does not display in the following instances:

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

You cannot clear the restart checkpoint for a baseline transfer.

Steps

1. In the **Protection** tab of the **Health/Volume** details page, right-click the relationship in the topology view for the data transfer you want to abort and select **Abort**.

The Abort Transfer dialog box is displayed.

- If you want to clear the restart checkpoint for a transfer that is not a baseline transfer, select Clear Checkpoints.
- 3. Click Continue.

The Abort Transfer dialog box is closed, and the status of the abort operation displays at the top of the Health/Volume details page along with a link to the job details.

- 4. Click the **View details** link to go to the **Protection/Job** details page for additional details and to view job progress.
- 5. Click each job task to view its details.
- 6. Click the Back arrow on your browser to return to the **Health/Volume** details page.

The abort operation is finished when all job tasks successfully complete.

Quiescing a protection relationship from the Protection/Volume Relationships page

From the Protection/Volume Relationships page, you can quiesce a protection relationship to temporarily prevent data transfers from occurring. You might quiesce a relationship when you want to create a Snapshot copy of a SnapMirror destination volume that contains a database, and you want to ensure that its contents are stable during the Snapshot copy operation.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

About this task

The quiesce action does not display in the following instances:

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

Steps

1. To quiesce transfers for one or more protection relationships, from the **Protection/Volume Relationships** page, select one or more volumes and, on the toolbar, click **Quiesce**.

The Quiesce dialog box is displayed.

2. Click Continue.

The status of the quiesce job is displayed at the top of the Health/Volume details page, along with a link to the job details.

- 3. Click the View details link to go to the Protection/Job details page for additional details and job progress.
- 4. Click the **Back** arrow on your browser to return to the **Protection/Volume Relationships** page.

The quiesce job is finished when all job tasks are successfully completed.

Quiescing a protection relationship from the Health/Volume details page

You can quiesce a protection relationship to temporarily prevent data transfers from occurring. You might quiesce a relationship when you want to create a Snapshot copy of a SnapMirror destination volume that contains a database, and you want to ensure that its contents are stable during the Snapshot copy.

Before you begin

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

You must have set up Workflow Automation.

About this task

The quiesce action does not display in the following instances:

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

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, right-click the relationship in the topology view for the protection relationship that you want to quiesce.
- Select Quiesce from the menu.
- 3. Click Yes to continue.

The status of the quiesce job is displayed at the top of the Health/Volume details page, along with a link to the job details.

- 4. Click the View details link to go to the Protection/Job details page for additional details and job progress.
- 5. Click the Back arrow on your browser to return to the **Health/Volume** details page.

The quiesce job is finished when all job tasks are successfully completed.

Breaking a SnapMirror relationship from the Protection/Volume Relationships page

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

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. From the **Protection/Volume Relationships** page, select one or more volumes with protection relationships for which you want to stop data transfers and, on the toolbar, click **Break**.

The Break Relationship dialog box is displayed.

- 2. Click **Continue** to break the relationship.
- In the Protection/Volume Relationships page, verify in the Relationship State column that the relationship is broken.

The Relationship State column is hidden by default, so you might need to select it in the show/hide column list [13].

Breaking a SnapMirror relationship from the Health/Volume details page

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

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, select from the topology the SnapMirror relationship you want to break.
- 2. Right-click the destination and select **Break** from the menu.

The Break Relationship dialog box is displayed.

- 3. Click **Continue** to break the relationship.
- 4. In the topology, verify that the relationship is broken.

Removing a protection relationship from the Protection/Volume Relationships page

From the Protection/Volume Relationships page, you can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. From the **Protection/Volume Relationships** page, select one or more volumes with protection relationships you want to remove and, on the toolbar, click **Remove**.

The Remove Relationship dialog box is displayed.

2. Click **Continue** to remove the relationship.

The relationship is removed from the Protection/Volume Relationships page.

Removing a protection relationship from the Health/Volume details page

You can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, select from the topology the SnapMirror relationship you want to remove.
- 2. Right-click the name of the destination and select **Remove** from the menu.

The Remove Relationship dialog box is displayed.

3. Click **Continue** to remove the relationship.

The relationship is removed from the Health/Volume details page.

Resuming scheduled transfers on a quiesced relationship from the Protection/Volume Relationships page

After you have quiesced a relationship to stop scheduled transfers from occurring, you can use **Resume** to reenable scheduled transfers so that data on the source or primary volume is protected. Transfers resume from a checkpoint, if one exists, at the next scheduled transfer interval.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

About this task

You can select no more than 10 quiesced relationships on which to resume transfers.

Steps

- 1. From the **Protection/Volume Relationships** page, select one or more volumes with quiesced relationships, and, on the toolbar, click **Resume**.
- 2. In the **Resume** dialog box, click **Continue**.

You are returned to the Protection/Volume Relationships page.

3. To view the related job tasks and to track their progress, click the job link that is displayed at the top of the

Protection/Volume Relationships page.

- 4. Do one of the following:
 - If only one job is displayed, in the Protection/Job details page click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
 - · If more than one job is displayed,
 - i. In the Protection/Jobs page, click the job for which you want to view the details.
 - ii. In the Protection/Job details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

 After the jobs finish, data transfers are resumed at the next scheduled transfer interval.

Resuming scheduled transfers on a quiesced relationship from the Health/Volume details page

After you have quiesced a relationship to stop scheduled transfers from occurring, you can use **Resume** on the Health/Volume details page to reenable scheduled transfers so that data on the source or primary volume is protected. Transfers resume from a checkpoint, if one exists, at the next scheduled transfer interval.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

Steps

1. In the **Protection** tab of the **Health/Volume** details page, right-click in the topology view a quiesced relationship that you want to resume.

Alternatively, select **Resume** from the **Actions > Relationship** menu.

2. In the Resume dialog box, click Continue.

You are returned to the Health/Volume details page.

- 3. To view the related job tasks and to track their progress, click the job link that is displayed at the top of the **Health/Volume** details page.
- 4. In the **Protection/Job** details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

After the jobs are complete, data transfers are resumed at the next scheduled transfer interval.

Initializing or updating protection relationships from the Protection/Volume Relationships page

From the Protection/Volume Relationships page, you can perform a first-time baseline transfer on a new protection relationship, or update a relationship if it is already initialized and you want to perform a manual, unscheduled incremental update to transfer immediately.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation.

Steps

1. From the **Protection/Volume Relationships** page, right-click a volume and select one or more volumes with relationships that you want to update or initialize, and then, on the toolbar, click **Initialize/Update**.

The Initialize/Update dialog box is displayed.

- 2. In the **Transfer Options** tab, select a transfer priority and the maximum transfer rate.
- 3. Click Source Snapshot Copies; then, in the Snapshot Copy column, click Default.

The Select Source Snapshot Copy dialog box is displayed.

- 4. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
- 5. Click Submit.

You are returned to the Initialize/Update dialog box.

- 6. If you selected more than one source to initialize or update, click **Default** for the next source for which you want to specify an existing Snapshot copy.
- 7. Click **Submit** to begin the initialization or update job.

The initialization or update job is started, you are returned to the Protection/Volume Relationships page, and a jobs link is displayed at the top of the page.

Click View Jobs on the Health/Volumes inventory page to track the status of each initialization or update job.

A filtered list of jobs is displayed.

- 9. Click each job to see its details.
- 10. Click the **Back** arrow on your browser to return to the **Protection/Volume Relationships** page.

The initialization or update operation is finished when all tasks successfully finish.

Initializing or updating protection relationships from the Health/Volume details page

You can perform a first-time baseline transfer on a new protection relationship, or update a relationship if it is already initialized and you want to perform a manual, unscheduled incremental update to transfer data immediately.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation.

Steps

- 1. From the **Protection** tab of the **Health/Volume** details page, locate in the topology the protection relationship that you want to initialize or update, and then right-click it.
- 2. Select Initialize/Update from the menu.

Alternatively, from the **Actions** menu, select **Relationship** > **Initialize/Update** to initialize or update the relationship for which you are currently viewing the details.

The Initialize/Update dialog box is displayed.

- 3. In the **Transfer Options** tab, select a transfer priority and the maximum transfer rate.
- 4. Click Source Snapshot Copies; then, in the Snapshot Copy column, click Default.

The Select Source Snapshot Copy dialog box is displayed.

- 5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
- 6. Click Submit.

You are returned to the Initialize/Update dialog box.

7. If you selected more than one source to initialize or update, click **Default** for the next read/write source for which you want to specify an existing Snapshot copy.

You cannot select a different Snapshot copy for data protection volumes.

8. Click **Submit** to begin the initialization or update job.

The initialization or update job is started, you are returned to the Health/Volume details page, and a jobs link is displayed at the top of the page.

9. Click View Jobs on the Health/Volume details page to track the status of each initialization or update job.

A filtered list of jobs is displayed.

- 10. Click each job to see its details.
- 11. Click the Back arrow on your browser to return to the **Health/Volume** details page.

The initialization or update operation is finished when all job tasks successfully complete.

Resynchronizing protection relationships from the Protection/Volume Relationships page

From the Protection/Volume Relationships page, you can resynchronize a relationship either to recover from an event that disabled your source volume or when you want to change the current source to a different volume.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

Steps

1. From the **Protection/Volume Relationships** page, select one or more volumes with quiesced relationships and, from the toolbar, click **Resynchronize**.

The Resynchronize dialog box is displayed.

- 2. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
- Click Source Snapshot Copies; then, in the Snapshot Copy column, click Default.

The Select Source Snapshot Copy dialog box is displayed.

- 4. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
- Click Submit.

You are returned to the Resynchronize dialog box.

- 6. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
- 7. Click **Submit** to begin the resynchronization job.

The resynchronization job is started, you are returned to the Protection/Volume Relationships page, and a jobs link is displayed at the top of the page.

8. Click **View Jobs** on the **Protection/Volume Relationships** page to track the status of each resynchronization job.

A filtered list of jobs is displayed.

9. Click the Back arrow on your browser to return to the Protection/Volume Relationships page.

The resynchronization operation is finished when all job tasks successfully finish.

Resynchronizing protection relationships from the Health/Volume details page

You can resynchronize data on a SnapMirror or SnapVault relationship that was broken and then the destination was made read/write so that data on the source matches the data on the destination. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation.

Steps

1. From the **Protection** tab of the **Health/Volume** details page, locate in the topology the protection relationship that you want to resynchronize and right-click it.

2. Select Resynchronize from the menu.

Alternatively, from the **Actions** menu, select **Relationship** > **Resynchronize** to resynchronize the relationship for which you are currently viewing the details.

The Resynchronize dialog box is displayed.

- 3. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
- 4. Click Source Snapshot Copies; then, in the Snapshot Copy column, click Default.

The Select Source Snapshot Copy dialog box is displayed.

- 5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
- 6. Click Submit.

You are returned to the Resynchronize dialog box.

- 7. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
- 8. Click **Submit** to begin the resynchronization job.

The resynchronization job is started, you are returned to the Health/Volume details page and a jobs link is displayed at the top of the page.

9. Click View Jobs on the Health/Volume details page to track the status of each resynchronization job.

A filtered list of jobs is displayed.

10. Click the Back arrow on your browser to return to the **Health/Volume** details page.

The resynchronization job is finished when all job tasks successfully complete.

Reversing protection relationships from the Protection/Volume Relationships 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 a read/write volume while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.

- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on reverse resynchronization relationships are the same as those on the original protection relationship.

If policies and schedules do not exist, they are created.

Steps

1. From the **Protection/Volume Relationships** page, select one or more volumes with relationships that you want to reverse, and, on the toolbar, click **Reverse Resync**.

The Reverse Resync dialog box is displayed.

2. Verify that the relationships displayed in the **Reverse Resync** dialog box are the ones for which you want to perform the reverse resynchronization operation, and then click **Submit**.

The reverse resynchronization operation is started, you are returned to the Protection/Volume Relationships page, and a jobs link is displayed at the top of the page.

3. Click **View Jobs** on the **Protection/Volume Relationships** page to track the status of each reverse resynchronization job.

A filtered list of jobs related to this operation is displayed.

4. Click the **Back** arrow on your browser to return to the **Protection/Volume Relationships** page.

The reverse resynchronization operation is finished when all job tasks successfully complete.

Reversing protection relationships from the Health/Volume details page

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.
- The relationship must not be a SnapVault relationship.
- · A protection relationship must already exist.
- The protection relationship must be broken.

- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on the reverse resynchronization relationship are the same as those on the original protection relationship.

If policies and schedules do not exist, they are created.

Steps

- 1. From the **Protection** tab of the **Health/Volume** details page, locate in the topology the SnapMirror relationship on which you want to reverse the source and destination, and right-click it.
- Select Reverse Resync from the menu.

The Reverse Resync dialog box is displayed.

3. Verify that the relationship displayed in the **Reverse Resync** dialog box is the one for which you want to perform the reverse resynchronization operation, and then click **Submit**.

The Reverse Resync dialog box is closed and a job link is displayed at the top of the Health/Volume details page.

 Click View Jobs on the Health/Volume details page to track the status of each reverse resynchronization job.

A filtered list of jobs is displayed.

5. Click the Back arrow on your browser to return to the **Health/Volume** details page.

The reverse resynchronization operation is finished when all job tasks are completed successfully.

Restoring data using the Health/Volumes inventory page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Health/Volumes inventory page.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot restore NTFS file streams.

The restore option is not available when:

• The volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered.

- The volume is a FlexGroup volume.
- The volume is configured for SnapMirror Synchronous replication.

Steps

- 1. In the **Health/Volumes** inventory page, select a volume from which you want to restore data.
- 2. From the toolbar, click **Restore**.

The Restore dialog box is displayed.

- 3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
- 4. Select the items you want to restore.

You can restore the entire volume, or you can specify folders and files you want to restore.

- Select the location to which you want the selected items restored; either Original Location or Alternate Location.
- Click Restore.

The restore process begins.

Restoring data using the Health/Volume details page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Health/Volume details page.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot restore NTFS file streams.

The restore option is not available when:

- The volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered.
- The volume is a FlexGroup volume.
- The volume is configured for SnapMirror Synchronous replication.

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, right-click in the topology view the name of the volume that you want to restore.
- 2. Select **Restore** from the menu.

Alternatively, select **Restore** from the **Actions** menu to protect the current volume for which you are viewing the details.

The Restore dialog box is displayed.

- 3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
- 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:
 - i. Select the cluster, SVM, and volume to which you want to restore.
 - ii. In the Name table, select a directory name.
 - iii. Click Select Directory.
- 7. Click Restore.

The restore process begins.



If a restore operation fails between Cloud Volumes ONTAP HA clusters with an NDMP error, you may need to add an explicit AWS route in the destination cluster so that the destination can communicate with the source system's cluster management LIF. You perform this configuration step using OnCommand Cloud Manager.

Creating resource pools

You can use the Create Resource Pool dialog box to group aggregates for provisioning purposes.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Resource pools can contain aggregates from different clusters, but the same aggregate cannot belong to different resource pools.

Steps

- 1. In the left navigation pane, click **Protection > Resource Pools**.
- 2. In the **Protection/Resource Pools** page, click **Create**.
- 3. Follow the instructions in the **Create Resource Pool** dialog box to provide a name and description and to add aggregates as members to the resource pool you want to create.

Editing resource pools

You can edit an existing resource pool when you want to change the resource pool name and the description.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The **Edit** button is enabled only when one resource pool is selected. If more than one resource pool is selected, the **Edit** button is disabled.

Steps

- 1. In the left navigation pane, click **Protection > Resource Pools**.
- 2. Select one resource pool from the list.
- 3. Click Edit.

The Edit Resource Pool window is displayed.

- 4. Edit the resource pool name and description as needed.
- 5. Click Save.

The new name and description are displayed in the resource pool list.

Viewing resource pools inventory

You can use the Protection/Resource Pools page to view the resource pool inventory and to monitor the remaining capacity for each resource pool.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Protection > Resource Pools**.

The resource pool inventory is displayed.

Adding resource pool members

A resource pool consists of a number of member aggregates. You can add aggregates to existing resource pools to increase the amount of space available for secondary volume provisioning.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can add no more than 200 aggregates to a resource pool at one time. Aggregates shown in the Aggregates dialog box do not belong to any other resource pool.

Steps

- 1. In the left navigation pane, click **Protection > Resource Pools**.
- 2. Select a resource pool from the **Resource Pools** list.

The resource pool members are displayed in the area below the resource pool list.

3. In the resource pool member area, click Add.

The Aggregates dialog box is displayed.

- 4. Select one or more aggregates.
- 5. Click Add.

The dialog box is closed and the aggregates are displayed in the member list for the selected resource pool.

Removing aggregates from resource pools

You can remove aggregates from an existing resource pool: for example, when you want to use an aggregate for some other purpose.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Resource pool members are displayed only when a resource pool is selected.

Steps

- 1. In the left navigation pane, click **Protection > Resource Pools**.
- 2. Select the resource pool from which you want to remove member aggregates.

The list of member aggregates is displayed in the Members pane.

3. Select one or more aggregates.

The **Remove** button is enabled.

4. Click Remove.

A warning dialog box is displayed.

Click Yes to continue.

The selected aggregates are removed from the Members pane.

Deleting resource pools

You can delete resource pools when they are no longer needed. For example, you might want to redistribute the member aggregates from one resource pool to several other resource pools, making the original resource pool obsolete.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The **Delete** button is enabled only when at least one resource pool is selected.

Steps

- 1. In the left navigation pane, click **Protection > Resource Pools**.
- 2. Select the resource pool you want to delete.
- 3. Click Delete.

The resource pool is removed from the resource pool list and its aggregates are removed from the members list.

Understanding SVM associations

storage virtual machine (SVM) associations are mappings from a source SVM to a destination SVM that are used by partner applications for resource selection and secondary volume provisioning.

Associations are always created between a source SVM and a destination SVM, regardless of whether the destination SVM is a secondary destination or a tertiary destination. You cannot use a secondary destination SVM as a source to create an association with a tertiary destination SVM.

You can associate SVMs in three ways:

Associate any SVM

You can create an association between any primary source SVM and one or more destination SVMs. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVMs. For example, you might want applications from several different sources at different locations to be backed up to one or more destination SVMs in one location.

· Associate a particular SVM

You can create an association between a specific source SVM and one or more specific destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one

another, you can choose this option to associate a specific source SVM to a specific destination SVM that is assigned to only that client.

· Associate with an external SVM

You can create an association between a source SVM and an external flexible volume of a destination SVM.

SVM and resource pool requirements to support storage services

You can better ensure conformance in partner applications if you observe some SVM association and resource pool requirements that are specific to storage services: for example, when you associate SVM and create resource pools in Unified Manager to support a protection topology in a storage service provided by a partner application.

Some applications partner with the Unified Manager server to provide services that automatically configure and execute SnapMirror or SnapVault backup protection between source volumes and protection volumes in secondary or tertiary locations. To support these protection storage services, you must use Unified Manager to configure the necessary SVM associations and resource pools.

To support storage service single-hop or cascaded protection, including replication from a SnapMirror source or SnapVault primary volume to either destination SnapMirror or to SnapVault backup volumes that reside in secondary or tertiary locations, observe the following requirements:

- SVM associations must be configured between the SVM containing the SnapMirror source or SnapVault primary volume and any SVM on which either a secondary volume or a tertiary volume resides.
 - For example, to support a protection topology in which source volume Vol_A resides on SVM_1, and SnapMirror secondary destination volume Vol_B resides on SVM_2, and tertiary SnapVault backup volume Vol_ C resides on SVM_3, you must use the Unified Manager web UI to configure a SnapMirror association between SVM_1 and SVM_2 and a SnapVault backup association between SVM_1 and SVM_3.

In this example, any SnapMirror association or SnapVault backup association between SVM_2 and SVM_3 is not necessary and is not used.

- To support a protection topology in which both source volume Vol_A and SnapMirror destination volume Vol_B reside on SVM_1, you must configure a SnapMirror association between SVM_1 and SVM_1.
- The resource pools must include cluster aggregate resources available to the associated SVMs.

You configure resource pools through the Unified Manager web UI and then assign through the partner application the storage service secondary target and tertiary target nodes.

Creating SVM associations

The Create Storage Virtual Machine Associations wizard enables partner protection applications to associate a source storage virtual machine (SVM) with a destination SVM for use with SnapMirror and SnapVault relationships. Partner applications use these associations at the time of initial provisioning of destination volumes to determine which resources to select.

Before you begin

- The SVM you are associating must already exist.
- You must have the OnCommand Administrator or Storage Administrator role.

About this task

For any source SVM and relationship type, you can choose only one destination SVM on each destination cluster.

Changing associations using the delete and create functions affects only future provisioning operations. It does not move existing destination volumes.

Steps

- 1. In the left navigation pane, click **Protection > SVM Associations**.
- In the Protection/Storage Virtual Machine Associations page, click Create.

The Create Storage Virtual Machine Associations wizard is launched.

3. Select one of the following sources:

• Any

Choose this option when you want to create an association between any primary SVM source to one or more destination SVM. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVM. For example, you might want applications from several different sources at different locations backed up to one or more destination SVM in one location.

Single

Choose this option when you want to select a specific source SVM associated with one or more destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, choose this option to associate a specific SVM source to a specific SVM destination that is assigned only to that client.

None (External)

Choose this option when you want to create an association between a source SVM and an external flexible volume of a destination SVM.

- 4. Select one or both of the protection relationship types you want to create:
 - SnapMirror
 - SnapVault
- 5. Click Next.
- 6. Select one or more SVM protection destination.
- 7. Click Finish.

Viewing SVM associations

You can use the Protection/Storage Virtual Machine Associations page to view existing

SVM associations and their properties and to determine if additional SVM associations are required.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Protection > SVM Associations**.

The list of SVM associations and their properties is displayed.

Deleting SVM associations

You can delete SVM associations for partner applications to remove the secondary provisioning relationship between source and destination SVMs; for example, you might do this when the destination SVM is full and you want to create a new SVM protection association.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The **Delete** button is disabled until at least one SVM association is selected. Changing associations using the delete and create functions affects only future provisioning operations; it does not move existing destination volumes.

Steps

- 1. In the left navigation pane, click **Protection > SVM Associations**.
- Select at least one SVM association.

The **Delete** button is enabled.

3. Click Delete.

A warning dialog box is displayed.

4. Click Yes to continue.

The selected SVM association is removed from the list.

What jobs are

A job is a series of tasks that you can monitor using Unified Manager. Viewing jobs and their associated tasks enables you to determine if a they have completed successfully.

Jobs are initiated when you create SnapMirror and SnapVault relationships, when you perform any relationship

operation (break, edit, quiesce, remove, resume, resynchronize, and reverse resync), when you perform data restoration tasks, when you log in to a cluster, and so on.

When you initiate a job, you can use the Protection/Jobs page and the Protection/Job details page to monitor the job and the progress of the associated job tasks.

Monitoring jobs

You can use the Protection/Jobs page to monitor job status and to view job properties such as storage service type, state, submitted time, and completed time to determine whether or not a job has successfully completed.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Protection > Jobs**.

The Protection/Jobs page is displayed.

- 2. View the **State** column to determine the status of those jobs currently running.
- 3. Click on a job name to view details about that particular job.

The Protection/Job details page is displayed.

Viewing job details

After you start a job, you can track its progress from the Protection/Job details page and monitor the associated tasks for possible errors.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Protection > Jobs**.
- 2. In the **Protection/Jobs** page, click a job name in the **Name** column to display the list of tasks associated with the job.
- 3. Click on a task to display additional information in the **Task Details** pane and the **Task Messages** pane to the right of the task list.

Aborting jobs

You can use the Protection/Jobs page to abort a job if it is taking too long to finish, is encountering too many errors, or is no longer needed. You can abort a job only if its status and type allow it. You can abort any running job.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Protection > Jobs**.
- 2. From the list of jobs, select one job, and then click **Abort**.
- 3. At the confirmation prompt, click **Yes** to abort the selected job.

Retrying a failed protection job

After you have taken measures to fix a failed protection job, you can use **Retry** to run the job again. Retrying a job creates a new job using the original job ID.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can retry only one failed job at a time. Selecting more than one job disables the **Retry** button. Only jobs of the type Protection Configuration and Protection Relationship Operation can be retried.

Steps

- 1. In the left navigation pane, click **Protection > Jobs**.
- 2. From the list of jobs, select a single failed Protection Configuration or Protection Relationship Operation type job.

The **Retry** button is enabled.

3. Click Retry.

The job is restarted.

Description of Protection relationships windows and dialog boxes

You can view and manage protection-related details such as resource pools, SVM associations, and protection jobs. You can use the appropriate Configuration/Health Thresholds page to configure global health threshold values for aggregates, volumes, and relationships.

Protection/Resource Pools page

The Protection/Resource Pools page displays existing resource pools and their members, and enables you to create, monitor, and manage resource pools for provisioning purposes.

Command buttons

The command buttons enable you to perform the following tasks:

Create

Launches the Create Resource Pool dialog box, which you can use to create resource pools.

• Edit

Enables you to edit the name and description of the resource pools that you create.

Delete

Enables you to delete one or more resource pools.

Resource Pools list

The Resource Pools list displays (in tabular format) the properties of existing resource pools.

Resource Pool

Displays the name of the resource pool.

Description

Describes the resource pool.

SnapLock Type

Displays the SnapLock type that is being used by the aggregates in the resource pool. Valid values for SnapLock type are Compliance, Enterprise, and Non-SnapLock. A resource pool can contain aggregates of only one SnapLock type.

Total Capacity

Displays the total capacity (in MB, GB, and so on) of the resource pool.

Used Capacity

Displays the amount of space (in MB, GB, and so on) that is used in the resource pool.

Available Capacity

Displays the amount of space (in MB, GB, and so on) that is available in the resource pool.

Used %

Displays the percentage of space that is used in the resource pool.

Members list command buttons

The Members list command buttons enable you to perform the following tasks:

Add

Enables you to add members to the resource pool.

Delete

Enables you to delete one or more members from the resource pool.

Members list

The Members list displays (in tabular format) the resource pool members and their properties when a resource pool is selected.

Status

Displays the current status of the member aggregate. The status can be Critical (\bigotimes), Error (\bigoplus), Warning (\bigwedge), or Normal (\bigotimes).

Aggregate Name

Displays the name of the member aggregate.

State

Displays the current state of the aggregate, which can be one of the following:

· Offline

Read or write access is not allowed.

· Online

Read and write access to the volumes that are hosted on this aggregate is allowed.

Restricted

Limited operations (such as parity reconstruction) are allowed, but data access is not allowed.

Creating

The aggregate is being created.

Destroying

The aggregate is being destroyed.

Failed

The aggregate cannot be brought online.

• Frozen

The aggregate is (temporarily) not serving requests.

· Inconsistent

The aggregate has been marked corrupted; you should contact technical support.

Iron Restricted

Diagnostic tools cannot be run on the aggregate.

Mounting

The aggregate is in the process of mounting.

Partial

At least one disk was found for the aggregate, but two or more disks are missing.

· Quiescing

The aggregate is being quiesced.

· Quiesced

The aggregate is quiesced.

Reverted

The revert of an aggregate is completed.

Unmounted

The aggregate has been unmounted.

Unmounting

The aggregate is being taken offline.

Unknown

The aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

By default, this column is hidden.

Cluster

Displays the name of the cluster to which the aggregate belongs.

Node

Displays the name of the node on which the aggregate resides.

Total Capacity

Displays the total capacity (in MB, GB, and so on) of the aggregate.

Used Capacity

Displays the amount of space (in MB, GB, and so on) that is used in the aggregate.

Available Capacity

Displays the amount of space (in MB, GB, and so on) that is available in the aggregate.

Used %

Displays the percentage of space that is used in the aggregate.

Disk Type

Displays the RAID configuration type, which can be one of the following:

- RAID0: All the RAID groups are of type RAID0.
- RAID4: All the RAID groups are of type RAID4.
- RAID-DP: All the RAID groups are of type RAID-DP.
- RAID-TEC: All the RAID groups are of type RAID-TEC.
- Mixed RAID: The aggregate contains RAID groups of different RAID types (RAID0, RAID4, RAID-DP, and RAID-TEC).

By default, this column is hidden.

Create Resource Pool dialog box

You can use the Create Resource Pool dialog box to name and describe a new resource pool and to add aggregates to and delete aggregates from that resource pool.

Resource Pool Name

The text boxes enable you to add the following information to create a resource pool:

Enables you to specify a resource pool name.

Description

Enables you to describe a resource pool.

Members

Displays the members of the resource pool. You can also add and delete members.

Command buttons

The command buttons enable you to perform the following tasks:

Add

Opens the Aggregates dialog box so that you can add aggregates from a specific cluster to the resource pool. You can add aggregates from different clusters, but the same aggregates cannot be added to more than one resource pool.

Remove

Enables you to remove selected aggregates from the resource pool.

Create

Creates the resource pool. This button is not enabled until information has been entered in the Resource Pool Name or Description fields.

Cancel

Discards the changes and closes the Create Resource Pool dialog box.

Edit Resource Pool dialog box

You can use the Edit Resource Pool dialog box to change the name and description of an existing resource pool. For example, if the original name and description is inaccurate or incorrect, you can change them so they are more precise.

Text boxes

The text boxes enable you to change the following information for the selected resource pool:

Resource Pool Name

Enables you to enter a new name.

Description

Enables you to enter a new description.

Command buttons

The command buttons enable you to perform the following tasks:

Save

Saves the changes to the resource pool name and description.

Cancel

Discards the changes and closes the Edit Resource Pool dialog box.

Aggregates dialog box

You can use the Aggregates dialog box to select the aggregates that you want to add to your resource pool.

Command buttons

The command buttons enable you to perform the following tasks:

Add

Adds the selected aggregates to the resource pool. The Add button is not enabled until at least one aggregate is selected.

Cancel

Discards the changes, and closes the Aggregates dialog box.

Aggregates list

The Aggregates list displays (in tabular format) the names and properties of monitored aggregates.

Status

Displays the current status of a volume. The status can be Critical (\bigotimes), Error (\bigoplus), Warning (\bigwedge), or Normal (\bigotimes).

You can move the pointer over the status to view more information about the event or events generated for the volume.

Aggregate Name

Displays the name of the aggregate.

State

Displays the current state of the aggregate, which can be one of the following:

Offline

Read or write access is not allowed.

Restricted

Limited operations (such as parity reconstruction) are allowed, but data access is not allowed.

Online

Read and write access to the volumes that are hosted on this aggregate is allowed.

Creating

The aggregate is being created.

Destroying

The aggregate is being destroyed.

Failed

The aggregate cannot be brought online.

Frozen

The aggregate is (temporarily) not serving requests.

Inconsistent

The aggregate has been marked corrupted; you should contact technical support.

· Iron Restricted

Diagnostic tools cannot be run on the aggregate.

Mounting

The aggregate is in the process of mounting.

Partial

At least one disk was found for the aggregate, but two or more disks are missing.

· Quiescing

The aggregate is being quiesced.

Quiesced

The aggregate is quiesced.

Reverted

The revert of an aggregate is completed.

Unmounted

The aggregate is offline.

Unmounting

The aggregate is being taken offline.

Unknown

The aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

Cluster

Displays the name of the cluster on which the aggregate resides.

Node

Displays the name of the storage controller that contains the aggregate.

Total Capacity

Displays the total data size (in MB, GB, and so on) of the aggregate. By default, this column is hidden.

Committed Capacity

Displays the total space (in MB, GB, and so on) that is committed for all the volumes in the aggregate. By default, this column is hidden.

Used Capacity

Displays the amount of space (in MB, GB, and so on) that is used in the aggregate.

· Available Capacity

Displays the amount of space (in MB, GB, and so on) that is available for data in the aggregate. By default, this column is hidden.

Available %

Displays the percentage of space that is available for data in the aggregate. By default, this column is hidden.

• Used %

Displays the percentage of space that is used by data in the aggregate.

RAID Type

Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, RAID-TEC, or Mixed RAID.

Protection/Storage Virtual Machine Associations page

The Protection/Storage Virtual Machine Associations page enables you to view existing SVM associations between source and destination SVMs and to create new SVM associations for use by partner applications to create SnapMirror and SnapVault relationships.

Command buttons

The command buttons enable you to perform the following tasks:

Create

Opens the Create Storage Virtual Machine Associations wizard.

Delete

Enables you to delete the selected SVM associations.

storage virtual machine (SVM) Associations list

The Storage Virtual Machine Associations list displays in a table the source and destination SVM associations that have been created and the type of protection relationship allowed for each association.

Source Storage Virtual Machine

Displays the name of the source SVM.

Source Cluster

Displays the name of the source cluster.

Destination Storage Virtual Machine

Displays the name of the destination SVM.

Destination Cluster

Displays the name of the destination cluster.

Type

Displays the type of protection relationship. Relationship types are either SnapMirror or SnapVault.

Create Storage Virtual Machine Associations wizard

The Create Storage Virtual Machine Associations wizard enables you to associate source and destination storage virtual machines (SVMs) for use in SnapMirror and SnapVault protection relationships.

Select Source SVM

The Select Source Storage Virtual Machine panel enables you to select the source, or primary, SVM in the SVM association.

Any

Enables you to create an association between any SVM source to one or more destination, or secondary, SVM. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVM. For example, you might want applications from several different sources at different locations backed up to one or more destination SVM in one location.

Single

Enables you to associate a specific source SVM with one or more destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, choose this option to associate a specific SVM source to a specific SVM destination that is assigned only to that client.

None (External)

Enables you to create an association between a source SVM and an external flexible volume of a destination SVM.

Storage Virtual Machine

Lists the names of the available source SVMs

Cluster

Lists the clusters on which each SVM resides

Allow these kinds of relationships

Enables you to select the relationship type for the association:

SnapMirror

Specifies a SnapMirror relationship as the association type. Selecting this option enables data replication from the selected sources to the selected destinations.

SnapVault

Specifies a SnapVault relationship as the association type. Selecting this option enables backups from the selected primary locations to the selected secondary locations.

Select Protection Destinations

The Select Protection Destinations panel of the Create Storage Virtual Machine Associations wizard enables you to select where to copy or replicate the data. You can create an association on only one destination SVM per cluster.

Command buttons

The command buttons enable you to perform the following tasks:

Next

Advances you to the next page in the wizard.

Back

Returns you to the previous page in the wizard.

Finish

Applies your selections and creates the association.

Cancel

Discards the selections and closes the Create Storage Virtual Machine Associations wizard.

Protection/Jobs page

The Protection/Jobs page enables you to view the current status and other information about all partner application protection jobs that are currently running, as well as jobs that have completed. You can use this information to see which jobs are still running and whether a job has succeeded or failed.

Command buttons

The command buttons enable you to perform the following tasks:

Abort

Aborts the selected job. This option is available only if the selected job is running.

Retry

Restarts a failed job of type Protection Configuration or Protection Relationship Operation. You can retry only one failed job at a time. If more than one failed job is selected, the **Retry** button is disabled. You cannot retry failed storage service jobs.

Refresh

Refreshes the list of jobs and the information associated with them.

Jobs list

The Jobs list displays, in tabular format, a list of the jobs that are in progress. By default, the list displays only the jobs generated within the past week. You can use column sorting and filtering to customize which jobs are displayed.

Status

Displays the current status of a job. The status can be Error ([1]) or Normal ([7]).

Job Id

Displays the identification number of the job. By default, this column is hidden.

The job identification number is unique and is assigned by the server when it starts the job. You can search for a particular job by entering the job identification number in the text box provided by the column filter.

Name

Displays the name of the job.

Type

Displays the job type. The job types are as follows:

Cluster Acquisition

A Workflow Automation job is rediscovering a cluster.

Protection Configuration

A protection job is initiating Workflow Automation workflows, such as cron schedules, SnapMirror policy creation, and so on.

Protection Relationship Operation

A protection job is running SnapMirror operations.

Protection Workflow Chain

A Workflow Automation job is executing multiple workflows.

Restore

A restore job is running.

Cleanup

The job is cleaning up storage service member artifacts that are no longer needed for restore purposes.

Conform

The job is checking the configuration of storage service members to ensure that they conform.

Destroy

The job is destroying a storage service.

Import

The job is importing unmanaged storage objects into an existing storage service.

Modify

The job is modifying attributes of an existing storage service.

Subscribe

The job is subscribing members to a storage service.

Unsubscribe

The job is unsubscribing members from a storage service.

Update

A protection update job is running.

WFA Configuration

A Workflow Automation job is pushing cluster credentials and synchronizing database caches.

State

Displays the running state of the job. State options are as follows:

Aborted

The job has been aborted.

Aborting

The job is in the process of aborting.

Completed

The job has finished.

Running

The job is running.

Submitted Time

Displays the time the job was submitted.

Duration

Displays the amount of time the job took to complete. This column is displayed by default.

Completed Time

Displays the time the job finished. By default, this column is hidden.

Protection/Job details page

The Protection/Job details page enables you to view status and other information about specific protection job tasks that are running, that are queued, or that have completed. You can use this information to monitor protection job progress and to troubleshoot job failures.

Job summary

The job summary displays the following information:

- Job ID
- Type
- State
- Submitted Time
- · Completed Time
- Duration

Command buttons

The command buttons enable you to perform the following tasks:

Refresh

Refreshes the task list and the properties associated with each task.

View Jobs

Returns you to the Protection/Jobs page.

Job tasks list

The Job tasks list displays in a table all the tasks associated with a specific job and the properties related to each task.

Started Time

Displays the day and time the task started. By default, the most recent tasks are displayed at the top of the column and older tasks are displayed at the bottom.

Type

Displays the type of task.

State

The state of a particular task:

Completed

The task has finished.

Queued

The task is about to run.

Running

The task is running.

Waiting

A job has been submitted and some associated tasks are waiting to be queued and executed.

Status

Displays the task status:

• Error ([]])

The task failed.

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.

Advanced Secondary Settings dialog box

You can use the Advanced Secondary Settings dialog box to enable version-flexible replication, multiple copy backup, and space-related settings on a secondary volume. You might use the Advanced Secondary Settings dialog box when you want to change enable or disable the current settings.

Space-related settings maximize the amount of data being stored, including the following: deduplication, data compression, autogrow, and space guarantee.

The dialog box includes the following fields:

Enable Version-Flexible Replication

Enables SnapMirror with version-flexible replication. Version-flexible replication enables SnapMirror protection of a source volume even if the destination volume is running under an earlier version of ONTAP than that of the source volume, as long as both source and destination are running ONTAP 8.3 or later.

Enable Backup

If version-flexible replication is enabled, also enables multiple Snapshot copies of the SnapMirror source data to be transferred to and retained at the SnapMirror destination.

Enable Deduplication

Enables deduplication on the secondary volume in a SnapVault relationship so that duplicate data blocks are eliminated to achieve space savings. You might use deduplication when space savings are at least 10 percent and when data overwrite rate is not rapid. Deduplication is often used for virtualized environments, file shares, and backup data. This setting is disabled by default. When enabled, this operation is initiated after each transfer.

Enable Compression

Enables transparent data compression. You might use compression when space savings are at least 10 percent, when the potential overhead is acceptable, and when there are sufficient system resources for compression to complete during nonpeak hours. In a SnapVault relationship, this setting is disabled by default. Compression is available only when deduplication is selected.

· Compress Inline

Enables immediate space savings by compressing data before writing data to disk. You might use inline compression when your system has no more than 50 percent utilization during peak hours, and when the system can accommodate new writes and additional CPU during peak hours. This setting is available only when "Enable Compression" is selected.

Enable Autogrow

Enables you to automatically grow the destination volume when the free space percentage is below the specified threshold, as long as space is available on the associated aggregate.

Maximum Size

Sets the maximum percentage to which a volume can grow. The default is 20 percent greater than the source volume size. A volume does not grow automatically if the current size is greater than or equal to the maximum autogrow percentage. This field is enabled only when the autogrow setting is enabled.

Increment Size

Specifies the percentage increment by which the volume automatically grows before reaching the maximum percentage of the source volume.

Space Guarantee

Ensures that enough space is allocated on the secondary volume so that data transfers always succeed. The space guarantee setting can be one of the following:

- File
- Volume
- ° None

For example, you might have a 200 GB volume that contains files totaling 50 GB; however, those files hold only 10 GB of data. Volume guarantee allocates 200 GB to the destination volume, regardless of contents on the source. File guarantee allocates 50 GB to ensure that enough space is reserved for files on the source; selecting None in this scenario means that only 10 GB is allocated on the destination for the actual space used by file data on the source.

The space guarantee is set to Volume by default.

Command buttons

The command buttons enable you to perform the following tasks:

Apply

Saves the selected efficiency settings and applies them when you click **Apply** in the Configure Protection dialog box.

Cancel

Discards your selections and closes the Advanced Destination Settings dialog box.

Related information

NetApp Technical Report 3966: NetApp Data Compression and Deduplication Deployment and Implementation Guide (Clustered Data ONTAP)

Advanced Destination Settings dialog box

You can use the Advanced Destination Settings dialog box to enable space guarantee settings on a destination volume. You might select advanced settings when space guarantee is disabled on the source, but you want it enabled on the destination. The settings for deduplication, compression, and autogrow in a SnapMirror relationship are inherited from the source volume and cannot be changed.

Space Guarantee

Ensures that enough space is allocated on the destination volume so that data transfers always succeed. The space guarantee setting can be one of the following:

• File

Space guarantee for files is not available in ONTAP 8.3.

- Volume
- None

For example, you might have a 200-GB volume that contains files totaling 50 GB; however, those files hold only 10 GB of data. Volume guarantee allocates 200 GB to the destination volume, regardless of contents on the source. File guarantee allocates 50 GB to ensure that enough space is reserved for source files on the destination; selecting **None** in this scenario means that only 10 GB is allocated on the destination for the actual space used by file data on the source.

The space guarantee is set to Volume by default.

Restore dialog box

You can use the Restore dialog box to restore data to a volume from a specific Snapshot copy.

Restore from

The Restore from area enables you to specify from where you want to restore data.

Volume

Specifies the volume from which you want to restore data. By default, the volume on which you initiated the restore action is selected. You can select a different volume from the drop-down list that contains all the volumes with protection relationships to the volume on which you initiated the restore action.

Snapshot copy

Specifies which Snapshot copy you want to use to restore data. By default, the most recent Snapshot copy is selected. You can also select a different Snapshot copy from the drop-down list. The Snapshot copy list changes according to which volume is selected.

· List maximum of 995 files and directories

By default a maximum of 995 objects are shown in the list. You can deselect this checkbox if you want to view all objects within the selected volume. This operation may take some time if the number of items is very large.

Select items to restore

The Select items to restore area enables you to select either the entire volume or specific files and folders you want to restore. You can select a maximum of 10 files, folders, or a combination of both. When the maximum number of items is selected, the item selection check boxes are disabled.

Path field

Displays the path to the data you want to restore. You can either navigate to the folder and files you want to restore, or you can type the path. This field is empty until you select or type a path. Clicking after you have chosen a path moves you up one level in the directory structure.

· Folders and files list

Displays the contents of the path you entered. By default, the root folder is initially displayed. Clicking a folder name displays the contents of the folder.

You can select items to restore as follows:

- When you enter the path with a particular file name specified in the path field, the specified file is displayed
 in the Folders and Files.
- When you enter a path without specifying a particular file, the contents of the folder are displayed in the Folders and Files list, and you can select up to 10 files, folders, or a combination of both to restore.

If a folder contains more than 995 items, a message displays to indicate there are too many items to display, and if you proceed with the operation all items in the specified folder are restored. You can deselect the "List maximum of 995 files and directories" checkbox if you want to view all objects within the selected volume.



You cannot restore NTFS file streams.

Restore to

The Restore to area enables you to specify where you want to restore the data.

Original Location in Volume Name

Restores the selected data to the directory on the source from which the data was originally backed up.

Alternate Location

Restores the selected data to a new location:

Restore Path

Specifies an alternate path for restoring the selected data. The path must already exist. You can use the **Browse** button to navigate to the location where you want the data restored, or you can enter the path manually using the format cluster://svm/volume/path.

Preserve directory hierarchy

When checked, preserves the structure of the original file or directory. For example, if the source is /A/B/C/myFile.txt and the destination is /X/Y/Z, Unified Manager restores the data using the following directory structure on the destination: /X/Y/Z/A/B/C/myFile.txt.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards your selections and closes the Restore dialog box.

Restore

Applies your selections and begins the restore process.

Browse Directories dialog box

You can use the Browse Directories dialog box when you want to restore data to a directory on a cluster and SVM that is different from the original source. The original source cluster and volume are selected by default.

The Browse Directories dialog box enables you to select the cluster, SVM, volume, and directory path to which you want data restored.

Cluster

Lists the available cluster destinations to which you can restore. By default, the cluster of the original source volume is selected.

SVM drop-down list

Lists the available SVM available for the selected cluster. By default, the SVM of the original source volume is selected.

Volume

Lists all of the read/write volumes in a selected SVM. You can filter the volumes by name and by space available. The volume with the most space is listed first, and so on, in descending order. By default, the original source volume is selected.

File path text box

Enables you to type the file path to which you want data restored. The path you enter must already exist.

Name

Displays the names of the available folders for the selected volume. Clicking a folder in the Name list displays the subfolders, if any exist. Files contained in the folders are not displayed. Clicking after you have selected a folder moves you up one level in the directory structure.

Command buttons

The command buttons enable you to perform the following tasks:

Select Directory

Applies your selections and closes the Browse Directories dialog box. If no directory is selected, this button is disabled.

Cancel

Discards your selections and closes the Browse Directories dialog box.

Configure Protection dialog box

You can use the Configure Protection dialog box to create SnapMirror and SnapVault relationships for all read, write, and data protection volumes on clusters to ensure that the data on a source volume or primary volume is replicated.

Source tab

Topology view

Displays a visual representation of the relationship that you are creating. The source in the topology is highlighted by default.

Source Information

Displays details about the selected source volumes, including the following information:

- Source cluster name
- Source SVM name
- Cumulative volume total size

Displays the total size of all the source volumes that are selected.

Cumulative volume used size

Displays the cumulative volume used size for all the selected source volumes.

Source volume

Displays the following information in a table :

Source Volume

Displays the names of the selected source volumes.

Type

Displays the volume type.

SnapLock Type

Displays the SnapLock type of the volume. The options are Compliance, Enterprise, and Non-SnapLock.

Snapshot Copy

Displays the Snapshot copy that is used for the baseline transfer. If the source volume is read/write, the value of Default in the Snapshot copy column indicates that a new Snapshot copy is created by default, and is used for the baseline transfer. If the source volume is a data protection volume, the value of Default in the Snapshot copy column indicates that no new Snapshot copy is created, and all existing Snapshot copies are transferred to the destination. Clicking the Snapshot copy value displays a list of Snapshot copies from which you can select an existing Snapshot copy to use for the baseline transfer. You cannot select a different default Snapshot copy if the source type is data protection.

SnapMirror tab

Enables you to specify a destination cluster, storage virtual machine (SVM), and aggregate for a protection relationship, as well as a naming convention for destinations while creating a SnapMirror relationship. You can also specify a SnapMirror policy and schedule.

Topology view

Displays a visual representation of the relationship that you are creating. The SnapMirror destination resource in the topology is highlighted by default.

Destination Information

Enables you to select the destination resources for a protection relationship:

Advanced link

Launches the Advanced Destination Settings dialog box when you are creating a SnapMirror relationship.

Cluster

Lists the clusters that are available as protection destination hosts. This field is required.

storage virtual machine (SVM)

Lists the SVMs that are available on the selected cluster. A cluster must be selected before the SVM list is populated. This field is required.

Aggregate

Lists the aggregates that are available on the selected SVM. A cluster must be selected before the Aggregate list is populated. This field is required. The Aggregate list displays the following information:

Rank

When multiple aggregates satisfy all the requirements for a destination, the rank indicates the priority in which the aggregate is listed, according to the following conditions:

- A. An aggregate that is located on a different node than the source volume node is preferred to enable fault domain separation.
- B. An aggregate on a node with fewer volumes is preferred to enable load balancing across nodes in a cluster.
- C. An aggregate that has more free space than other aggregates is preferred to enable capacity balancing.

A rank of 1 means that the aggregate is the most preferred according to the three criteria.

Aggregate Name

Name of the aggregate

- Available Capacity
- Amount of space that is available on the aggregate for data
- Resource Pool

Name of the resource pool to which the aggregate belongs

Naming Convention

Specifies the default naming convention that is applied to the destination volume. You can accept the naming convention that is provided, or you can create a custom one. The naming convention can have the following attributes: %C, %M, %V, and %N, where %C is the cluster name, %M is the SVM name, %V is the source volume, and %N is the topology destination node name.

The naming convention field is highlighted in red if your entry is invalid. Clicking the "Preview Name" link displays a preview of the naming convention that you entered, and the preview text updates dynamically as you type a naming convention in the text field. A suffix between 001 and 999 is appended to the destination name when the relationship is created, replacing the nnn that displays in the preview text, with 001 being assigned first, 002 assigned second, and so on.

Relationship Settings

Enables you to specify the maximum transfer rate, SnapMirror policy, and schedule that the protection relationship uses:

Max Transfer Rate

Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored.

SnapMirror Policy

Specifies the ONTAP SnapMirror policy for the relationship. The default is DPDefault.

Create Policy

Launches the Create SnapMirror Policy dialog box, which enables you to create and use a new SnapMirror policy.

SnapMirror Schedule

Specifies the ONTAP SnapMirror policy for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.

· Create Schedule

Launches the Create Schedule dialog box, which enables you to create a new SnapMirror schedule.

SnapVault tab

Enables you to specify a secondary cluster, SVM, and aggregate for a protection relationship, as well as a naming convention for secondary volumes while creating a SnapVault relationship. You can also specify a SnapVault policy and schedule.

Topology view

Displays a visual representation of the relationship that you are creating. The SnapVault secondary resource in the topology is highlighted by default.

Secondary Information

Enables you to select the secondary resources for a protection relationship:

Advanced link

Launches the Advanced Secondary Settings dialog box.

Cluster

Lists the clusters that are available as secondary protection hosts. This field is required.

storage virtual machine (SVM)

Lists the SVMs that are available on the selected cluster. A cluster must be selected before the SVM list is populated. This field is required.

Aggregate

Lists the aggregates that are available on the selected SVM. A cluster must be selected before the Aggregate list is populated. This field is required. The Aggregate list displays the following information:

Rank

When multiple aggregates satisfy all the requirements for a destination, the rank indicates the priority in which the aggregate is listed, according to the following conditions:

- A. An aggregate that is located on a different node than the primary volume node is preferred to enable fault domain separation.
- B. An aggregate on a node with fewer volumes is preferred to enable load balancing across nodes in a cluster.
- C. An aggregate that has more free space than other aggregates is preferred to enable capacity balancing.

A rank of 1 means that the aggregate is the most preferred according to the three criteria.

Aggregate Name

Name of the aggregate

- Available Capacity
- Amount of space that is available on the aggregate for data
- Resource Pool

Name of the resource pool to which the aggregate belongs

Naming Convention

Specifies the default naming convention that is applied to the secondary volume. You can accept the naming convention that is provided, or you can create a custom one. The naming convention can have the following attributes: %C, %M, %V, and %N, where %C is the cluster name, %M is the SVM name, %V is the source volume, and %N is the topology secondary node name.

The naming convention field is highlighted in red if your entry is invalid. Clicking the "Preview Name" link displays a preview of the naming convention that you entered, and the preview text updates dynamically as you type a naming convention in the text field. If you type an invalid value, the invalid information displays as red question marks in the preview area. A suffix between 001 and 999 is appended to the secondary name when the relationship is created, replacing the nnn that displays in the preview text, with 001 being assigned first, 002 assigned second, and so on.

Relationship Settings

Enables you to specify the maximum transfer rate, SnapVault policy, and SnapVault schedule that the protection relationship uses:

Max Transfer Rate

Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored.

SnapVault Policy

Specifies the ONTAP SnapVault policy for the relationship. The default is XDPDefault.

Create Policy

Launches the Create SnapVault Policy dialog box, which enables you to create and use a new SnapVault policy.

SnapVault Schedule

Specifies the ONTAP SnapVault schedule for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.

· Create Schedule

Launches the Create Schedule dialog box, which enables you to create a SnapVault schedule.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards your selections, and closes the Configure Protection dialog box.

Apply

Applies your selections, and begins the protection process.

Create Schedule dialog box

The Create Schedule dialog box enables you to create a basic or advanced protection

schedule for SnapMirror and SnapVault relationship transfers. You might create a new schedule to increase the frequency of data transfers due to frequent data updates, or you might create a less frequent schedule when data changes infrequently.

Schedules cannot be configured for SnapMirror Synchronous relationships.

Destination Cluster

The name of the cluster you selected in the SnapVault tab or SnapMirror tab of the Configure Protection dialog box.

· Schedule Name

The name you provide for the schedule. Schedule names can consist of the characters A through Z, a through z, 0 through 9, as well as any of the following special characters: $! @ # $ % ^ & * () _ -.$ Schedule names may not include the following characters: < >.

· Basic or Advanced

The schedule mode you want to use.

Basic mode includes the following elements:

Repeat

How often a scheduled transfer occurs. Choices include hourly, daily, and weekly.

Day

When a repeat of weekly is selected, the day of the week a transfer occurs.

Time

When Daily or Weekly is selected, the time of day a transfer occurs.

Advanced mode includes the following elements:

Months

A comma-separated numerical list representing the months of the year. Valid values are 0 through 11, with zero representing January, and so on. This element is optional. Leaving the field blank implies that transfers occur every month.

Days

A comma-separated numerical list representing the day of the month. Valid values are 1 through 31. This element is optional. Leaving the field blank implies that a transfer occurs every day of the month.

Weekdays

A comma-separated numerical list representing the days of the week. Valid values are 0 through 6, with 0 representing Sunday, and so on. This element is optional. Leaving the field blank implies that a transfer occurs every day of the week. If a day of the week is specified but a day of the month is not specified, a transfer occurs only on the specified day of the week and not every day.

Hours

A comma-separated numerical list representing the number of hours in a day. Valid values are 0 through 23, with 0 representing midnight. This element is optional.

Minutes

A comma-separated numerical list representing the minutes in an hour. Valid values are 0 through 59. This element is required.

Create SnapMirror Policy dialog box

The Create SnapMirror Policy dialog box enables you to create a policy to set the priority for SnapMirror transfers. You use policies to maximize the efficiency of transfers from the source to the destination.

Destination Cluster

The name of the cluster you selected in the SnapMirror tab of the Configure Protection dialog box.

Destination SVM

The name of the SVM you selected in the SnapMirror tab of the Configure Protection dialog box.

Policy Name

The name you provide for the new policy. Policy names can consist of the characters A through Z, a through Z, 0 through 9, period (.), hyphen (-), and underscore ().

Transfer Priority

The priority at which a transfer runs for asynchronous operations. You can select either Normal or Low. Transfer relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority.

Comment

An optional field in which you can add comments about the policy.

Transfer Restart

Indicates what restart action to take when a transfer is interrupted by an abort operation or any type of failure, such as a network outage. You can select one of the following:

Always

Specifies that a new Snapshot copy is created before restarting a transfer, then, if one exists, the transfer is restarted from a checkpoint, followed by an incremental transfer from the newly created Snapshot copy..

Never

Specifies that interrupted transfers are never restarted.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards the selections and closes the Configure Protection dialog box.

Apply

Applies your selections and begins the protection process.

Create SnapVault Policy dialog box

The Create SnapVault Policy dialog box enables you to create a policy to set the priority for SnapVault transfers. You use policies to maximize the efficiency of transfers from the primary to the secondary volume.

Destination Cluster

The name of the cluster that you selected in the SnapVault tab of the Configure Protection dialog box.

Destination SVM

The name of the SVM that you selected in the SnapVault tab of the Configure Protection dialog box.

Policy Name

The name you provide for the new policy. Policy names can consist of the characters A through Z, a through Z, 0 through 9, period (.), hyphen (-), and underscore ().

Transfer Priority

The priority at which the transfer is run. You can select either Normal or Low. Transfer relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority. The default setting is Normal.

Comment

An optional field in which you can a add comment of up to 255 characters about the SnapVault policy.

Ignore Access Time

Specifies whether incremental transfers are ignored for files that have only their access time changed.

Replication Label

Lists in a table the rules associated with Snapshot copies selected by ONTAP that have a specific replication label in a policy. The following information and actions are also available:

Command buttons

The command buttons enable you to perform the following actions:

Add

Enables you to create a Snapshot copy label and retention count.

Edit Retention Count

Enables you to change the retention count for an existing Snapshot copy label. The retention count must be a number between 1 and 251. The sum of all retention counts for all rules cannot exceed 251.

Delete

Enables you to delete an existing Snapshot copy label.

Snapshot Copy Label

Displays the Snapshot copy label. If you select one or more volumes with the same local Snapshot copy policy, an entry for each label in the policy is displayed. If you select multiple volumes that have two or more local Snapshot copy policies, the table displays all labels from all policies

Schedule

Displays the schedule associated with each Snapshot copy label. If a label has more than one schedule associated with it, the schedules for that label are displayed in a comma-separated list. If you select multiple volumes with the same label but with different schedules, the schedule displays "Various" to indicate that more than one schedule is associated with the selected volumes.

Destination Retention Count

Displays the number of Snapshot copies with the specified label that are retained on the SnapVault secondary. Retention counts for labels with multiple schedules displays the sum of retention counts of each label and schedule pair. If you select multiple volumes with two or more local Snapshot copy policies, the retention count is empty.

Edit Relationship dialog box

You can edit an existing protection relationship to change the maximum transfer rate, the protection policy, or the protection schedule.

Destination Information

Destination Cluster

The name of the selected destination cluster.

Destination SVM

The name of the selected SVM

Relationship Settings

Enables you to specify the maximum transfer rate, SnapMirror policy, and schedule that the protection relationship uses:

Max Transfer Rate

Specifies the maximum rate at which baseline data is transferred between clusters over the network.

When selected, network bandwidth is limited to the value you specify. You can enter a numerical value and then select either kilobytes per second (KBps), megabytes per second (MBps), gigabytes per second (GBps), or terabytes per second (TBps). The maximum transfer rate that you specify must be greater than 1 KBps and less than 4 TBps. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. If the primary cluster and the secondary cluster are the same, this setting is disabled.

SnapMirror Policy

Specifies the ONTAP SnapMirror policy for the relationship. The default is DPDefault.

Create Policy

Launches the Create SnapMirror Policy dialog box, which enables you to create and use a new SnapMirror policy.

SnapMirror Schedule

Specifies the ONTAP SnapMirror policy for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.

· Create Schedule

Launches the Create Schedule dialog box, which enables you to create a new SnapMirror schedule.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards the selections and closes the Configure Protection dialog box.

Submit

Applies your selections and closes the Edit Relationship dialog box.

Initialize/Update dialog box

The Initialize/Update dialog box enables you to perform a first-time baseline transfer on a new protection relationship, or to update a relationship if it is already initialized and you want to perform a manual, unscheduled, incremental update.

Transfer Options tab

The Transfer Options tab enables you to change the initialization priority of a transfer and to change the bandwidth used during transfers.

Transfer Priority

The priority at which the transfer is run. You can select either Normal or Low. Relationships with policies that specify a normal transfer priority run before those that specify a low transfer priority. Normal is selected

by default.

Max Transfer Rate

Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is ignored. If you select more than one relationship with different maximum transfer rates, you can specify one of the following maximum transfer rate settings:

Use values specified during individual relationship setup or edit

When selected, initialization and update operations use the maximum transfer rate specified at the time of each relationship's creation or edit. This field is available only when multiple relationships with different transfer rates are being initialized or updated.

Unlimited

Indicates that there is no bandwidth limitation on transfers between relationships. This field is available only when multiple relationships with different transfer rates are being initialized or updated.

Limit bandwidth to

When selected, network bandwidth is limited to the value you specify. You can enter a numerical value and then select either kilobytes per second (KBps), Megabytes per second (MBps), Gigabytes per second (GBps), or Terabytes per second (TBps). The maximum transfer rate that you specify must be greater than 1 KBps and less than 4 TBps.

Source Snapshot Copies tab

The Source Snapshot Copies tab displays the following information about the source Snapshot copy that is used for the baseline transfer:

Source Volume

Displays the names of the corresponding source volumes.

Destination Volume

Displays the names of the selected destination volumes.

Source Type

Displays the volume type. The type can be either Read/write or Data Protection.

Snapshot Copy

Displays the Snapshot copy that is used for the data transfer. Clicking the Snapshot copy value displays the Select Source Snapshot Copy dialog box, in which you can select a specific Snapshot copy for your transfer, depending on the type of protection relationship that you have and the operation that you are performing. The option to specify a different Snapshot copy is not available for data protection type sources.

Command buttons

The command buttons enable you to perform the following tasks:

Cancel

Discards your selections and closes the Initialize/Update dialog box.

Submit

Saves your selections and starts the initialize or update job.

Resynchronize dialog box

The Resynchronize dialog box enables you to resynchronize data on a SnapMirror or SnapVault relationship that was previously broken and then the destination was made a read/write volume. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Resynchronization Options tab

The Resynchronization Options tab enables you to set the transfer priority and the maximum transfer rate for the protection relationship that you are resynchronizing.

Transfer Priority

The priority at which the transfer is run. You can select either Normal or Low. Relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority.

Max Transfer Rate

Specifies the maximum rate at which data is transferred between clusters over the network. When selected, network bandwidth is limited to the value that you specify. You can enter a numerical value and then select either kilobytes per second (KBps), megabytes per second (MBps), gigabytes per second (GBps), or TBps. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. However, if you are running ONTAP 8.2, and the primary cluster and the secondary cluster are the same, this setting is disabled.

Source Snapshot Copies tab

The Source Snapshot Copies tab displays the following information about the source Snapshot copy that is used for the baseline transfer:

Source Volume

Displays the names of the corresponding source volumes.

Destination Volume

Displays the names of the selected destination volumes.

Source Type

Displays the volume type: either Read/write or Data Protection.

Snapshot Copy

Displays the Snapshot copy that is used for the data transfer. Clicking the Snapshot copy value displays the Select Source Snapshot Copy dialog box, in which can select a specific Snapshot copy for your transfer, depending on the type of protection relationship you have and the operation you are performing.

Command buttons

Submit

Begins the resynchronization process and closes the Resynchronize dialog box.

Cancel

Cancels your selections and closes the Resynchronize dialog box.

Select Source Snapshot Copy dialog box

You use the Select Source Snapshot Copy dialog box to select a specific Snapshot copy to transfer data between protection relationships, or you select the default behavior, which varies depending on whether you are initializing, updating, or resynchronizing a relationship, and whether the relationship is a SnapMirror or SnapVault.

Default

Enables you to select the default behavior for determining which Snapshot copy is used for initialize, update, and resynchronize transfers for SnapVault and SnapMirror relationships.

If you are performing a SnapVault transfer, the default behavior for each operation is as follows:

Operation	Default SnapVault behavior when source is read/write	Default SnapVault behavior when source is Data Protection (DP)
Initialize	Creates a new Snapshot copy and transfers it.	Transfers the last exported Snapshot copy.
Update	Transfers only labeled Snapshot copies, as specified in the policy.	Transfers the last exported Snapshot copy.
Resynchronize	Transfers all labeled Snapshot copies created after the newest common Snapshot copy.	Transfers the newest labeled Snapshot copy.

If you are performing a SnapMirror transfer, the default behavior for each operation is as follows:

Operation	Default SnapMirror behavior	Default SnapMirror behavior when relationship is second hop in a SnapMirror to SnapMirror cascade
Initialize	Creates a new Snapshot copy and transfers it and all Snapshot copies created prior to the new Snapshot copy.	Transfers all Snapshot copies from the source.
Update	Creates a new Snapshot copy and transfers it and all Snapshot copies created prior to the new Snapshot copy.	Transfers all Snapshot copies.
Resynchronize	Creates a new Snapshot copy and then transfers all Snapshot copies from the source.	Transfers all Snapshot copies from the secondary volume to the tertiary volume, and deletes any data added after creation of the newest common Snapshot copy.

Existing Snapshot Copy

Enables you to select an existing Snapshot copy from the list if Snapshot copy selection is allowed for that operation.

Snapshot Copy

Displays the existing Snapshot copies from which you can select for a transfer.

Date Created

Displays the date and time the Snapshot copy was created. Snapshot copies are listed from most recent to least recent, with the most recent at the top of the list.

If you are performing a SnapVault transfer and you want to select an existing Snapshot copy to transfer from a source to a destination, the behavior for each operation is as follows:

Operation	SnapVault behavior when specifying a Snapshot copy	SnapVault behavior when specifying a Snapshot copy in a cascade
Initialize	Transfers the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Update	Transfers the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

Operation	SnapVault behavior when specifying a Snapshot copy	SnapVault behavior when specifying a Snapshot copy in a cascade
Resynchronize	Transfers the selected Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

If you are performing a SnapMirror transfer and you want to select an existing Snapshot copy to transfer from a source to a destination, the behavior for each operation is as follows:

Operation	SnapMirror behavior when specifying a Snapshot copy	SnapMirror behavior when specifying a Snapshot copy in a cascade
Initialize	Transfers all Snapshot copies on the source, up to the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Update	Transfers all Snapshot copies on the source, up to the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Resynchronize	Transfers all Snapshot copies from the source, up to the selected Snapshot copy, and then deletes any data added after creation of the newest common Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

Command buttons

The command buttons enable you to perform the following tasks:

Submit

Submits your selections and closes the Select Source Snapshot Copy dialog box.

Cancel

Discards your selections and closes the Select Source Snapshot Copy dialog box.

Reverse Resync dialog box

When you have a protection relationship that is broken because the source volume is disabled and the destination is made a read/write volume, reverse resynchronization enables you to reverse the direction of the relationship so that the destination becomes the new source and the source becomes the new destination.

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write, while you repair or replace the source, update the source, and reestablish the relationship. When you perform a reverse resync operation, data on the source that is

newer than the data on the common Snapshot copy is deleted.

Before reverse resync

Displays the source and destination of a relationship before a reverse resync operation.

Source Volume

The name and location of the source volume before a reverse resync operation.

Destination Volume

The name and location of the destination volume before a reverse resync operation.

After reverse resync

Displays what the source and destination of a relationship is after a reserve resync operation.

Source Volume

The name and location of the source volume after a reverse resync operation.

Destination Volume

The name and location of the destination volume after a reverse resync operation.

Command buttons

The command buttons enable you to perform the following actions:

Submit

Begins the reverse resynchronization process.

Cancel

Closes the Reverse Resync dialog box without initiating a reverse resync operation.

Protection/Volume Relationships page

The Protection/Volume Relationships page displays information about protection relationships on the storage system.

Use the **Export** button to export the details of all the relationships to a comma-separated values (.csv) file.

Relationship Status

Displays the current status of the protection relationship.

The status can be one of Error (1), Warning (Λ) , or Normal (2).

Lag Status

Displays the lag status for managed relationships, and for unmanaged relationships that have a schedule

associated with that relationship. Lag status can be:

• Error (

The lag duration is greater than or equal to the lag error threshold.

Warning (<u>A</u>)

The lag duration is greater than or equal to the lag warning threshold.

∘ Normal (**⊘**)

The lag duration is within normal limits.

Not Applicable

The lag status is not applicable for synchronous relationships because a schedule cannot be configured.

Transfer Status

Displays the transfer status for the protection relationship. The transfer status can be one of the following:

Aborting

SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

· Checking

The destination volume is undergoing a diagnostic check and no transfer is in progress.

Finalizing

SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

· Idle

Transfers are enabled and no transfer is in progress.

∘ In-Sync

The data in the two volumes in the synchronous relationship are synchronized.

· Out-of-Sync

The data in the destination volume is not synchronized with the source volume.

Preparing

SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.

· Queued

SnapMirror transfers are enabled. No transfers are in progress.

Quiesced

SnapMirror transfers are disabled. No transfer is in progress.

· Quiescing

A SnapMirror transfer is in progress. Additional transfers are disabled.

Transferring

SnapMirror transfers are enabled and a transfer is in progress.

Transitioning

The asynchronous transfer of data from the source to the destination volume is complete, and the transition to synchronous operation has started.

Waiting

A SnapMirror transfer has been initiated, but some associated tasks are waiting to be queued.

Relationship Type

Displays the relationship type used to replicate a volume. Relationship types include:

- Asynchronous Mirror
- Asynchronous Vault
- StrictSync
- Sync

Source SVM

Displays the name of the source SVM.

If the message Resource-key not discovered is displayed, this might indicate that the SVM exists on the cluster but has not yet been added to the Unified Manager inventory, or that the SVM was created after the cluster's last refresh. You must ensure that the SVM exists, or perform a rediscovery on the cluster to refresh the list of resources.

You can move your pointer over the source SVM to view information such as the cluster, volume type, protocols allowed, and spaced used. You can view more details about the SVM by clicking on the SVM name.

Source Volume

Displays the source volume being protected. You can view more details about the source volume by clicking the source volume name.

If the message Resource-key not discovered is displayed, this might indicate that the volume exists on the cluster but has not yet been added to the Unified Manager inventory, or that the volume was created after the cluster's last refresh. You must ensure that the volume exists, or perform a rediscovery on the cluster to refresh the list of resources.

Destination SVM

Displays the name of the destination SVM.

You can move your pointer over the destination SVM to view information such as the cluster, volume type, protocols allowed, and space used. You can view more details about the SVM by clicking on the SVM name.

Destination Volume

Displays the name of the destination volume.

You can move the pointer over a volume to view information such as the aggregate containing the volume, qtree quota overcommitted space, status of the last volume move operation, and space allocated in the volume. You can also view the details of related objects such as the SVM to which the volume belongs, the aggregate to which the volume belongs, and all the volumes that belong to this aggregate.

Lag Duration

Displays the amount of time that the data on the mirror lags behind the source.

The lag duration should be close to, or equal to, 0 seconds for StrictSync relationships.

Last Successful Update

Displays the time of the last successful SnapMirror or SnapVault operation.

The last successful update is not applicable for synchronous relationships.

Last Transfer Duration

Displays the time taken for the last data transfer to complete.

The transfer duration is not applicable for StrictSync relationships because the transfer should be simultaneous.

Last Transfer Size

Displays the size, in bytes, of the last data transfer.

The transfer size is not applicable for StrictSync relationships.

Relationship Health

Displays the relationship heath of the cluster.

Relationship State

Displays the the mirror state of the SnapMirror relationship.

Unhealthy Reason

The reason the relationship is in an unhealthy state.

Source Cluster

Displays the name of the source cluster for the SnapMirror relationship.

Source Node

Displays the name of the source node for the SnapMirror relationship.

Destination Cluster

Displays the name of the destination cluster for the SnapMirror relationship.

Destination Node

Displays the name of the destination node for the SnapMirror relationship.

Transfer Priority

Displays the priority at which a transfer runs. The transfer priority is Normal or Low. Normal priority transfers are scheduled before low priority transfers.

The transfer priority is not applicable for synchronous relationships because all transfers are treated with the same priority.

Policy

Displays the protection policy for the volume. You can click the policy name to view details associated with that policy, including the following information:

Transfer priority

Specifies the priority at which a transfer runs for asynchronous operations. The transfer priority is Normal or Low. Normal priority transfers are scheduled before low priority transfers. The default is Normal.

Ignore access time

Applies only to SnapVault relationships. This specifies whether incremental transfers ignore files which have only their access time changed. The values are either True or False. The default is False.

When Relationship is Out of Sync

Specifies the action ONTAP performs when a synchronous relationship is not able to be synchronized. StrictSync relationships will restrict access to the primary volume if there is a failure to synchronize with the secondary volume. Sync relationships do not restrict access to the primary if there is a failure to synchronize with the secondary.

· Tries limit

Specifies the maximum number of times to attempt each manual or scheduled transfer for a SnapMirror relationship. The default is 8.

Comments

Provides a text field for comments for specific to the selected policy.

SnapMirror label

Specifies the SnapMirror label for the first schedule associated with the Snapshot copy policy. The SnapMirror label is used by the SnapVault subsystem when you back up Snapshot copies to a SnapVault destination.

Retention settings

Specifies how long backups are kept, based on the time or the number of backups.

Actual Snapshot copies

Specifies the number of Snapshot copies on this volume that match the specified label.

Preserve Snapshot copies

Specifies the number of SnapVault Snapshot copies that are not deleted automatically even if the maximum limit for the policy is reached. The values are either True or False. The default is False.

Retention warning threshold

Specifies the Snapshot copy limit at which a warning is sent to indicate that the maximum retention limit is nearly reached.

Schedule

Displays the name of the protection schedule assigned to the relationship. You can click the schedule name to view details about the schedule.

The schedule is not applicable for synchronous relationships.

Version Flexible Replication

Displays either Yes, Yes with backup option, or None.

Protection/Volume Transfer Status (Historical) page

The Protection/Volume Transfer Status (Historical) page enables you to analyze the volume transfer trends over a period of time. This page also displays whether the volume transfer was a success or a failure.

Use the **Export** button to export the details of all the monitored volumes to a comma-separated values (.csv) file.



This page displays information for volumes in asynchronous relationships only - volumes in synchronous relationships are not shown.

Source Cluster

Displays the source cluster name.

Source SVM

Displays the storage virtual machine (SVM) name.

Source Volume

Displays the source volume name.

Destination Cluster

Displays the destination cluster name.

Destination SVM

Displays the destination SVM name.

Destination Volume

Displays the destination volume name.

Operation Result

Displays whether volume transfer was successful.

Transfer Start Time

Displays the volume transfer start time.

Transfer End Time

Displays the volume transfer end time.

Transfer Duration

Displays the time taken (in hours) to complete the volume transfer.

Transfer Size

Displays the size (in MB) of the transferred volume.

Operation Type

Displays the type of volume transfer.

Protection/Volume Transfer Rate (Historical) page

The Protection/Volume Transfer Rate (Historical) page enables you to analyze the amount of data volume that is transferred on a day-to-day basis. This page also provides details about daily volume transfers and the time required to complete the transfer operation.

Use the **Export** button to export the details of all the monitored volumes to a comma-separated values (.csv) file.



This page displays information for volumes in asynchronous relationships only - volumes in synchronous relationships are not shown.

Total Transfer Size

Displays the total size of the volume transfer in gigabytes.

Day

Displays the day on which the volume transfer was initiated.

End Time

Displays the volume transfer end time with date.

Executing protection workflows using OnCommand Workflow Automation

You can integrate OnCommand Workflow Automation with Unified Manager to execute workflows for your storage classes and monitor SVMs with Infinite Volume that do not have storage classes.

Configuring a connection between Workflow Automation and Unified Manager

You can configure a secure connection between OnCommand Workflow Automation (WFA) and Unified Manager. Connecting to Workflow Automation enables you to use protection features such as SnapMirror and SnapVault configuration workflows, as well as commands for managing SnapMirror relationships.

Before you begin

- The installed version of Workflow Automation must be 4.2 or greater.
- You must have installed "WFA pack for managing Clustered Data ONTAP" version 9.5.0 or greater on the WFA server. You can download the required pack from the NetAppStorage Automation Store.

WFA pack for managing ONTAP

 You must have the name of the database user that you created in Unified Manager to support WFA and Unified Manager connections.

This database user must have been assigned the Integration Schema user role.

- You must be assigned either the Administrator role or the Architect role in Workflow Automation.
- You must have the host address, port number 443, user name, and password for the Workflow Automation setup.
- You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click **Workflow Automation** in the left Setup menu.
- In the OnCommand Unified Manager Database User area of the Setup/Workflow Automation page, select the name, and enter the password for the database user that you created to support Unified Manager and Workflow Automation connections.
- 3. In the **OnCommand Workflow Automation Credentials** area of the **Setup/Workflow Automation** page, enter the host name or IP address (IPv4 or IPv6), and the user name and password for the Workflow Automation setup.

You must use the Unified Manager server port (port 443).

- 4. Click Save.
- 5. If you use a self-signed certificate, click **Yes** to authorize the security certificate.

The Setup/Workflow Automation page displays.

6. Click **Yes** to reload the web UI, and add the Workflow Automation features.

Removing OnCommand Workflow Automation setup from Unified Manager

You can remove the OnCommand Workflow Automation setup from Unified Manager when you no longer want to use Workflow Automation.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click **Workflow Automation** in the left Setup menu.
- 2. In the Setup/Workflow Automation page, click Remove Setup.

What happens when OnCommand Workflow Automation is reinstalled or upgraded

Before reinstalling or upgrading OnCommand Workflow Automation, you must first remove the connection between OnCommand Workflow Automation and Unified Manager, and ensure that all OnCommand Workflow Automation currently running or scheduled jobs are stopped.

You must also manually delete Unified Manager from OnCommand Workflow Automation.

After you reinstall or upgrade OnCommand Workflow Automation, you must set up the connection with Unified Manager again.

Description of OnCommand Workflow Automation setup windows and dialog boxes

You can set up OnCommand Workflow Automation in Unified Manager by using the Setup/Workflow Automation page.

Setup/Workflow Automation page

The Setup/Workflow Automation page enables you to configure the settings to integrate OnCommand Workflow Automation with Unified Manager. You can also add, modify, or delete the settings.

You must have the OnCommand Administrator or Storage Administrator role.

Unified Manager Database User

This area enables you to enter the credentials of a database user that is required for pairing Unified Manager with Workflow Automation:

Name

Enables you to specify the user name of a database user that can be used to access data in the Unified Manager database. By default, no database user is selected. You can select a database user from the drop-down list.

Password

Enables you to specify a password for the specified user name.

OnCommand Workflow Automation Credentials

This area enables you to enter the credentials of an Workflow Automation account that is required for pairing with Unified Manager:

Hostname or IP Address

Specifies the name or IP address of the Workflow Automation host server, which is used to pair with Unified Manager.

Port

Displays the required the port number of the Workflow Automation host server, which is 443.

Username

Enables you to specify a user name that can be used to log in to Workflow Automation.

Password

Enables you to specify a password for the specified user name.

Command buttons

The command buttons enable you to remove, save, or cancel the setup options:

Remove Setup

Removes the Workflow Automation setup from Unified Manager.

Save

Saves the configuration settings for the selected option.

Managing performance using performance capacity and available IOPS information

Performance capacity indicates how much throughput you can get out of a resource

without surpassing the useful performance of that resource. When viewed using existing performance counters, performance capacity is the point at which you get the maximum utilization from a node or aggregate before latency becomes an issue.

Unified Manager collects performance capacity statistics from nodes and aggregates in each cluster. Performance capacity used is the percentage of performance capacity that is currently being used, and performance capacity free is the percentage of performance capacity that is still available.

While performance capacity free provides a percentage of the resource that is still available, available IOPS tells you the number of IOPS that can be added to the resource before reaching the maximum performance capacity. By using this metric, you can be sure that you can add workloads of a predetermined number of IOPS to a resource.

Monitoring the performance capacity information has the following benefits:

- Assists with workflow provisioning and balancing.
- Helps you prevent overloading a node or pushing its resources beyond the optimal point, thus reducing the need to troubleshoot.
- · Helps you determine with greater precision where additional storage equipment might be needed.

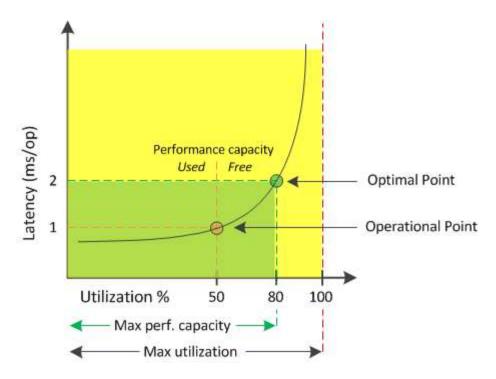
What performance capacity used is

The performance capacity used counter helps you to identify whether the performance of a node or an aggregate is reaching a point where the performance might degrade if the workloads increase. It can also show you if a node or aggregate is currently being overused during specific periods of time. Performance capacity used is similar to utilization, but the former provides more insight about the available performance capabilities in a physical resource for a specific workload.



Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

The optimal used performance capacity is the point at which a node or an aggregate has optimal utilization and latency (response time), and is being used efficiently. A sample latency versus utilization curve is shown for an aggregate in the following figure.



In this example, the *operational point* identifies that the aggregate is currently operating at 50% utilization with latency of 1.0 ms/op. Based on the statistics captured from the aggregate, Unified Manager determines that additional performance capacity is available for this aggregate. In this example, the *optimal point* is identified as the point when the aggregate is at 80% utilization with latency of 2.0 ms/op. Therefore, you can add more volumes and LUNs to this aggregate so that your systems are used more efficiently.

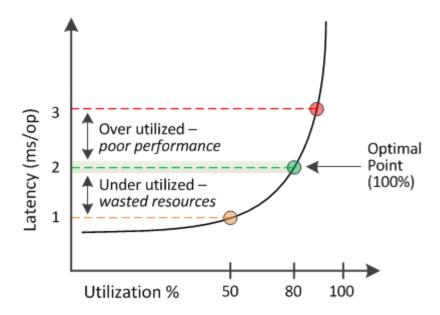
The performance capacity used counter is expected to be a larger number than the "utilization" counter because performance capacity adds in the impact on latency. For example, if a node or aggregate is 70% used, the performance capacity value may be in the 80% to 100% range, depending on the latency value.

In some cases, however, the utilization counter may be higher on the Dashboards/Performance page. This is normal because the dashboard refreshes the current counter values at each collection period; it does not display averages over a period of time like the other pages in the Unified Manager user interface. The performance capacity used counter is best used as an indicator of performance averaged over a period of time, whereas the utilization counter is best used for determining the instantaneous usage of a resource.

What the performance capacity used value means

The performance capacity used value helps you identify the nodes and aggregates that are currently being overutilized or underutilized. This enables you to redistribute workloads in order to make your storage resources more efficient.

The following figure shows the latency versus utilization curve for a resource and identifies, with colored dots, three areas where the current operational point could be located.



• A performance capacity used percentage equal to 100 is at the optimal point.

Resources are being used efficiently at this point.

• A performance capacity used percentage above 100 indicates that the node or aggregate is overutilized, and that workloads are receiving sub-optimal performance.

No new workloads should be added to the resource, and the existing workloads may need to be redistributed.

• A performance capacity used percentage below 100 indicates that the node or aggregate is underutilized, and that resources are not being used effectively.

More workloads can be added to the resource.



Unlike utilization, the performance capacity used percentage can be above 100%. There is no maximum percentage, but resources will typically be in the 110% to 140% range when they are being overutilized. Higher percentages would indicate a resource with serious issues.

What available IOPS is

The available IOPS counter identifies the remaining number of IOPS that can be added to a node or an aggregate before the resource reaches its limit. The total IOPS that a node can provide is based on the physical characteristics of the node—for example, the number of CPUs, the CPU speed, and the amount of RAM. The total IOPS that an aggregate can provide is based on the physical properties of the disks—for example, a SATA, SAS, or SSD disk.

While the performance capacity free counter provides the percentage of a resource that is still available, the available IOPS counter tells you the exact number of IOPS (workloads) can be added to a resource before reaching the maximum performance capacity.

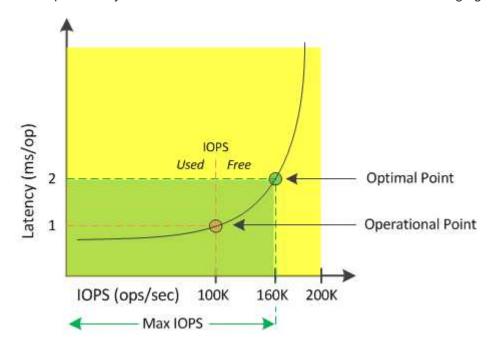
For example, if you are using a pair of FAS2520 and FAS8060 storage systems, a performance capacity free value of 30% means that you have some free performance capacity. However, that value does not provide visibility into how many more workloads you can deploy to those nodes. The available IOPS counter may show

that you have 500 available IOPS on the FAS8060, but only 100 available IOPS on the FAS2520.



Available IOPS data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

A sample latency versus IOPS curve for a node is shown in the following figure.



The maximum number of IOPS that a resource can provide is the number of IOPS when the performance capacity used counter is at 100% (the optimal point). The operational point identifies that the node is currently operating at 100K IOPS with latency of 1.0 ms/op. Based on the statistics captured from the node, Unified Manager determines that the maximum IOPS for the node is 160K, which means that there are 60K free or available IOPS. Therefore, you can add more workloads to this node so that your systems are used more efficiently.

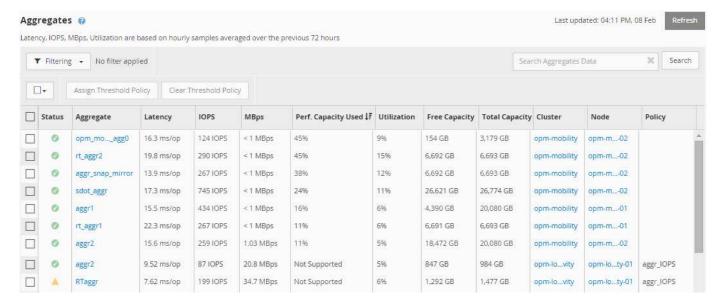


When there is minimal user activity in the resource, the available IOPS value is calculated assuming a generic workload based on approximately 4,500 IOPS per CPU core. This is because Unified Manager lacks the data to accurately estimate the characteristics of the workload being served.

Viewing node and aggregate performance capacity used values

You can monitor the performance capacity used values for all nodes or for all aggregates in a cluster, or you can view details for a single node or aggregate.

Performance capacity used values appear in the Performance Dashboard, Performance Inventory pages, Top Performers page, Create Threshold Policy page, Performance Explorer pages, and in detail charts. For example, the Performance/Aggregate Inventory page provides a column Perf. Capacity Used to view the performance capacity used value for all aggregates.



The status "N/A" is displayed when nodes are not installed with ONTAP 9.0 or later software.

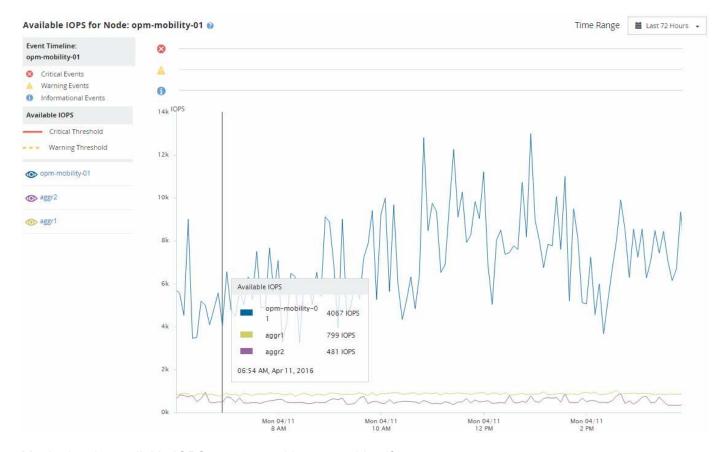
Monitoring the performance capacity used counter enables you to identify the following:

- · Whether any nodes or aggregates on any clusters have a high performance capacity used value
- · Whether any nodes or aggregates on any clusters have active performance capacity used events
- The nodes and aggregates that have the highest and lowest performance capacity used value in a cluster
- Latency and utilization counter values in conjunction with nodes or aggregates that have high performance capacity used values
- · How the performance capacity used values for nodes in an HA pair will be affected if one of the nodes fails
- · The busiest volumes and LUNs on an aggregate that has a high performance capacity used value

Viewing node and aggregate available IOPS values

You can monitor the available IOPS values for all nodes or for all aggregates in a cluster, or you can view details for a single node or aggregate.

Available IOPS values appear in the Performance Explorer page charts. For example, when viewing a node in the Performance/Node Explorer page, you can select the "Available IOPS" counter chart from the list so you can compare the available IOPS values for multiple aggregates on that node.



Monitoring the available IOPS counter enables you to identify:

- The nodes or aggregates that have the greatest available IOPS values to help determine where future workloads can be deployed.
- The nodes or aggregates that have the smallest available IOPS values to identify the resources you should monitor for potential future performance issues.
- The busiest volumes and LUNs on an aggregate that has a small available IOPS value.

Viewing performance capacity counter charts to identify issues

You can view performance capacity used charts for nodes and aggregates on the Performance Explorer page. This enables you to view detailed performance capacity data for the selected nodes and aggregates for a specific timeframe.

About this task

The standard counter chart displays the performance capacity used values for the selected nodes or aggregates. The Breakdown counter chart displays the total performance capacity values for the root object separated into usage based on user protocols versus background system processes. Additionally, the amount of free performance capacity is also shown.

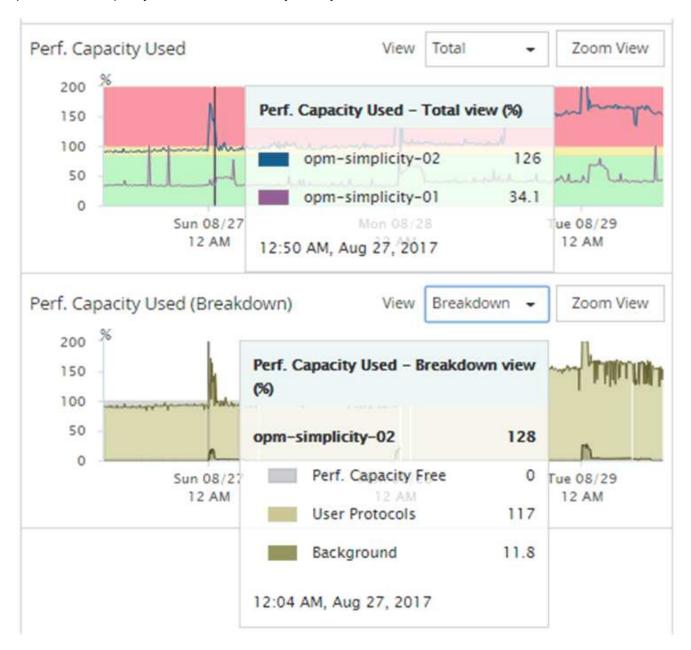


Because some background activities associated with system and data management are identified as user workloads and categorized as user protocols, the user protocols percentage may appear artificially high when those processes run. These processes typically run around midnight when cluster usage is low. If you see a spike in user protocol activity around midnight, verify if cluster backup jobs or other background activities are configured to run at that time.

Steps

- 1. Select the **Explorer** tab from a node or aggregate **Landing** page.
- 2. In the Counter Charts pane, click Choose charts, and then select the Perf. Capacity Used chart.
- 3. Scroll down until you can view the chart.

The colors of the standard chart show when the object is in the optimal range (yellow), when the object is underutilized (green), and when the object is overutilized (red). The Breakdown chart shows detailed performance capacity details for the root object only.



4. If you want to view either chart in a full size format, click **Zoom View**.

In this manner you can open multiple counter charts in a separate windows to compare performance capacity used values with IOPS or MBps values over the same timeframe.

Performance capacity used performance threshold conditions

You can create user-defined performance threshold policies so that events are triggered when the performance capacity used value for a node or aggregate exceeds the defined performance capacity used threshold setting.

Additionally, nodes can be configured with a "Performance capacity used takeover" threshold policy. This threshold policy totals the performance capacity used statistics for both nodes in an HA pair to determine whether either node would lack sufficient capacity if the other node fails. Because the workload during failover is the combination of the two partner nodes' workloads, the same performance capacity used takeover policy can be applied to both nodes.



This performance capacity used equivalency is generally true between nodes. However, if there is significantly more cross-node traffic destined for one of the nodes through its failover partner, the total performance capacity used when running all workloads on one partner node versus the other partner node could be slightly different depending on which node has failed.

The performance capacity used conditions can also be used as secondary performance threshold settings to create a combination threshold policy when defining thresholds for LUNs and volumes. The performance capacity used condition is applied to the aggregate or node on which the volume or LUN resides. For example, you can create a combination threshold policy using the following criteria:

Storage object	Performance counter	Warning threshold	Critical threshold	Duration
Volume	Latency	15 ms/op	25 ms/op	20 minutes

Combination threshold policies cause an event to be generated only when both conditions are breached for the entire duration.

Using the performance capacity used counter to manage performance

Typically, organizations want to operate with a performance capacity used percentage below 100 so that resources are being efficiently used while reserving some additional performance capacity to support peak period demands. You can use threshold policies to customize when alerts are sent for high performance capacity used values.

You can establish specific goals based on your performance requirements. For example, financial services firms might reserve more performance capacity to guarantee the timely execution of trades. These companies might want to set performance capacity used thresholds in the 70-80 percent range. Manufacturing companies with smaller margins might choose to reserve less performance capacity if they are willing to risk performance to better manage IT costs. These companies might set performance capacity used thresholds in the 85-95 percent range.

When the performance capacity used value exceeds the percentage set in a user-defined threshold policy, Unified Manager sends an alert email and adds the event to the Event Inventory page. This enables you to manage potential problems before they impact performance. These events can also be used as indicators that you need to make workload moves and changes within your nodes and aggregates.

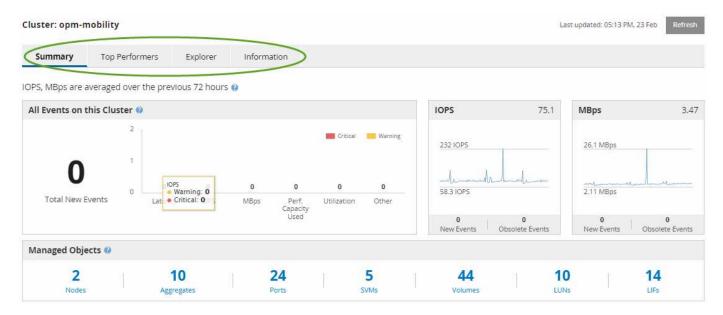
Monitoring cluster performance from the Performance Cluster Landing page

The Performance Cluster Landing page displays the high-level performance status of a selected cluster that is being monitored by an instance of Unified Manager. This page enables you to assess the overall performance of a specific cluster, and to quickly note, locate, or assign for resolution any cluster-specific events that are identified.

Understanding the Performance Cluster Landing page

The Performance Cluster Landing page provides a high-level performance overview of a selected cluster, with an emphasis on the performance status of the top 10 objects within the cluster. Performance issues are displayed at the top of the page, in the All Events on this Cluster panel.

The Performance Cluster Landing page provides a high-level overview of each cluster that is managed by an instance of Unified Manager. This page provides you with information about events and performance, and enables you to monitor and troubleshoot clusters. The following image shows an example of the Performance Cluster Landing page for the cluster called opm-mobility:



The event count on the Cluster Summary page may not match the event count on the Performance Event Inventory page. This is because the Cluster Summary page can show one event each in the Latency and Utilization bars when combination threshold policies have been breached, whereas the Performance Event Inventory page shows only one event when a combination policy has been breached.



If a cluster was removed from being managed by Unified Manager, the status **Removed** is displayed at the right of the cluster name at the top of the page.

Performance Cluster Landing page

The Performance Cluster Landing page displays the high-level performance status of a selected cluster. The page enables you to access complete details of each performance

counter for the storage objects on the selected cluster.

You can click the **Favorites** button () to add this object to your list of favorite storage objects. A blue button () indicates that this object is already a favorite.

The Performance Cluster Landing page includes four tabs that separate the cluster details into four areas of information:

- · Summary page
 - Cluster Events pane
 - Managed Objects pane
- Top Performers page
- Explorer page
- · Information page

Performance Cluster Summary page

The Performance Cluster Summary page provides a summary of the active events, IOPS performance, and MBps performance for a cluster. This page also includes the total count of the storage objects in the cluster.

Cluster performance events pane

The Cluster performance events pane displays performance statistics and all active events for the cluster. This is most helpful when monitoring your clusters and all cluster-related performance and events.

All Events on this Cluster pane

The All Events on this Cluster pane displays all active cluster performance events for the preceding 72 hours. The Total Active Events is displayed at the far left; this number represents the total of all New and Acknowledged events for all storage objects in this cluster. You can click the Total Active Events link to navigate to the Events Inventory page, which is filtered to display these events.

The Total Active Events bar graph for the cluster displays the total number of active critical and warning events:

- Latency (total for nodes, aggregates, SVMs, volumes, LUNs, and namespaces)
- IOPS (total for clusters, nodes, aggregates, SVMs, volumes, LUNs, and namespaces)
- MBps (total for clusters, nodes, aggregates, SVMs, volumes, LUNs, namespaces, ports, and LIFs)
- Performance Capacity Used (total for nodes and aggregates)
- Utilization (total for nodes, aggregates, and ports)
- Other (cache miss ratio for volumes)

The list contains active performance events triggered from user-defined threshold policies, system-defined threshold policies, and dynamic thresholds.

Graph data (vertical counter bars) is displayed in red () for critical events, and yellow () for warning events. Position your cursor over each vertical counter bar to view the actual type and number of events. You can click **Refresh** to update the counter panel data.

You can show or hide critical and warning events in the Total Active Events performance graph by clicking the **Critical** and **Warning** icons in the legend. If you hide certain event types, the legend icons are displayed in gray.

Counter panels

The counter panels display cluster activity and performance events for the preceding 72 hours, and includes the following counters:

IOPS counter panel

IOPS indicates the operating speed of the cluster in number of input/output operations per second. This counter panel provides a high-level overview of the cluster's IOPS health for the preceding 72-hour period. You can position your cursor over the graph trend line to view the IOPS value for a specific time.

· MBps counter panel

MBps indicates how much data has been transferred to and from the cluster in megabytes per second. This counter panel provides a high-level overview of the cluster's MBps health for the preceding 72-hour period. You can position your cursor over the graph trend line to view the MBps value for a specific time.

The number at the top right of the chart in the gray bar is the average value from the last 72-hour period. Numbers shown at the bottom and top of the trend line graph are the minimum and maximum values for the last 72-hour period. The gray bar below the chart contains the count of active (new and acknowledged) events and obsolete events from the last 72-hour period.

The counter panels contain two types of events:

Active

Indicates that the performance event is currently active (new or acknowledged). The issue causing the event has not corrected itself or has not been resolved. The performance counter for the storage object remains above the performance threshold.

Obsolete

Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the storage object is no longer above the performance threshold.

For **Active Events**, if there is one event, you can position your cursor over the event icon and click the event number to link to the appropriate Event Details page. If there is more than one event, you can click **View all Events** to display the Events Inventory page, which is filtered to show all events for the selected object counter type.

Managed Objects pane

The Managed Objects pane in the Performance Summary tab provides a top-level overview of the storage object types and counts for the cluster. This pane enables you to track the status of the objects in each cluster.

The managed objects count is point-in-time data as of the last collection period. New objects are discovered at 15-minute intervals.

Clicking the linked number for any object type displays the object performance inventory page for that object

type. The object inventory page is filtered to show only the objects on this cluster.

The managed objects are:

Nodes

A physical system in a cluster.

Aggregates

A set of multiple redundant array of independent disks (RAID) groups that can be managed as a single unit for protection and provisioning.

Ports

A physical connection point on nodes that is used to connect to other devices on a network.

SVMs

A virtual machine providing network access through unique network addresses. An SVM might serve data out of a distinct namespace, and is separately administrable from the rest of the cluster.

Volumes

A logical entity holding accessible user data through one or more of the supported access protocols. The count includes both FlexVol volumes and FlexGroup volumes; it does not include FlexGroup constituents or Infinite Volumes.

• LUNs

The identifier of a Fibre Channel (FC) logical unit or an iSCSI logical unit. A logical unit typically corresponds to a storage volume, and is represented within a computer operating system as a device.

• LIFs

A logical network interface representing a network access point to a node. The count includes all LIF types.

Top Performers page

The Top Performers page displays the storage objects that have the highest performance or the lowest performance, based on the performance counter you select. For example, in the SVMs category, you can display the SVMs that have the highest IOPS, or the highest latency, or the lowest MBps. This page is also shows if any of the top performers have any active performance events (New or Acknowledged).

The Top Performers page displays a maximum of 10 of each object. Note that the Volume object includes both FlexVol volumes and FlexGroup volumes; it does not include FlexGroup constituents or Infinite Volumes.

Time Range

You can select a time range for viewing the top performers; the selected time range applies to all storage objects. Available time ranges:

Last Hour

- Last 24 Hours
- Last 72 Hours (default)
- Last 7 Days

Metric

Click the **Metric** menu to select a different counter. Counter options are unique to the object type. For example, available counters for the **Volumes** object are **Latency**, **IOPS**, and **MBps**. Changing the counter reloads the panel data with the top performers based on the selected counter.

Available counters:

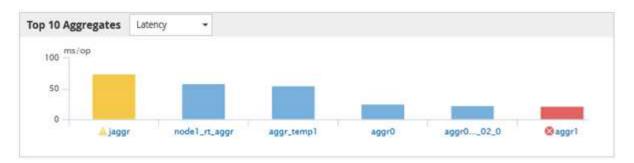
- Latency
- IOPS
- MBps
- Performance Capacity Used (for nodes and aggregates)
- Utilization (for nodes and aggregates)

• Sort

Click the **Sort** menu to select an ascending or descending sort for the selected object and counter. The options are **Highest to lowest** and **Lowest to highest**. These options enable you to view the objects with the highest performance or the lowest performance.

· Counter bar

The counter bar in the graph shows the performance statistics for each object, represented as a bar for that item. The bar graphs are color-coded. If the counter is not breaching a performance threshold, the counter bar is displayed in blue. If a threshold breach is active (a new or acknowledged event), the bar is displayed in the color for the event: warning events are displayed in yellow (), and critical events are displayed in red (). Threshold breaches are further indicated by severity event indicator icons for warning and critical events.



For each graph, the X axis displays the top performers for the selected object type. The Y axis displays units applicable to the selected counter. Clicking the object name link below each vertical bar graph element navigates to the Performance Landing page for the selected object.

Severity Event indicator

The **Severity Event** indicator icon is displayed at the left of an object name for active critical (**S**) or warning (**A**) events in the top performers graphs. Click the **Severity Event** indicator icon to view:

One event

Navigates to the Event details page for that event.

Two or more events

Navigates to the Event inventory page, which is filtered to display all events for the selected object.

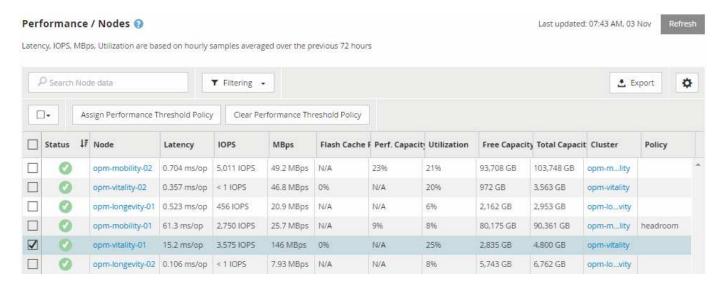
Export button

Creates a .csv file that contains the data that appears in the counter bar. You can choose to create the file for the single cluster you are viewing or for all clusters in the data center.

Monitoring performance using the Performance Inventory pages

The object inventory performance pages display performance information, performance events, and object health for all objects within an object type category. This provides you with an at-a-glance overview of the performance status of each object within a cluster, for example, for all nodes or all volumes.

Object inventory performance pages provide a high-level overview of object status, enabling you to assess the overall performance of all objects and compare object performance data. You can refine the content of object inventory pages by searching, sorting, and filtering. This is beneficial when monitoring and managing object performance, because it enables you to quickly locate objects with performance issues and to begin the troubleshooting process.



By default, objects on the performance inventory pages are sorted based on object performance criticality. Objects with new critical performance events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. All performance data is based on a 72-hour average.

You can easily navigate from the object inventory performance page to an object details page by clicking the object name in the object name column. For example, on the Performance/Nodes inventory page, you would click a node object in the **Nodes** column. The object details page provides in-depth information and detail about the selected object, including side-by-side comparison of active events.

Object monitoring using the Performance object inventory pages

The Performance object inventory pages enable you to monitor object performance based on the values of specific performance counters or based on performance events. This is beneficial because identifying objects with performance events enables you to investigate the cause of cluster performance issues.

The Performance object inventory pages display the associated counters, associated objects, and performance threshold policies for all objects in all clusters. These pages also enable you to apply performance threshold policies to objects. You can sort the page based on any column, and you can search across all object names or data.

You can export data from these pages to a comma-separated values (.csv) file by using the **Export** button, and then use the exported data to build reports.

Refining Performance inventory page contents

The inventory pages for performance objects contain tools to help you refine object inventory data content, enabling you to locate specific data quickly and easily.

Information contained within the Performance object inventory pages can be extensive, often spanning multiple pages. This kind of comprehensive data is excellent for monitoring, tracking, and improving performance; however, locating specific data requires tools to enable you to quickly locate the data for which you are looking. Therefore, the Performance object inventory pages contain functionality for searching, sorting, and filtering. Additionally, searching and filtering can work together to further narrow your results.

Searching on Object Inventory Performance pages

You can search strings on Object Inventory Performance pages. Use the **Search** field located at the top right of the page to quickly locate data based on either object name or policy name. This enables you to quickly locate specific objects and their associated data, or to quickly locate policies and view associated policy object data.

Steps

1. Perform one of the following options, based on your search requirements:

To locate this	Type this
A specific object	The object name into the Search field, and click Search . The object for which you searched and its related data is displayed.
A user-defined performance threshold policy	All or part of the policy name into the Search field, and click Search . The objects assigned to the policy for which you searched are displayed.

Sorting on the Object Inventory Performance pages

You can sort all data on Object Inventory Performance pages by any column in ascending

or descending order. This enables you to quickly locate object inventory data, which is helpful when examining performance or beginning a troubleshooting process.

About this task

The selected column for sorting is indicated by a highlighted column heading name and an arrow icon indicating the sorting direction at the right of the name. An up arrow indicates ascending order; a down arrow indicates descending order. The default sort order is by **Status** (event criticality) in descending order, with the most critical performance events listed first.

Steps

1. You can click a column name to toggle the sort order of the column in ascending or descending order.

The Object Inventory Performance page contents are sorted in ascending or descending order, based on the selected column.

Filtering data in the Object Inventory Performance pages

You can filter data in the Object Inventory Performance pages to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Object Inventory Performance pages to show only the results you have specified. This provides a very efficient method of displaying only the performance data in which you are interested.

About this task

You can use the Filtering panel to customize the grid view based on your preferences. Available filter options are based on the correlated object type being viewed in the grid. If filters are currently applied, an asterisk (*) displays at the left of the Filtering control.

Four types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are contains and starts with .
Number	The operators are greater than and less than .
Resource	The operators are name contains and name starts with.
Status	The operators are is and is not .

All three fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

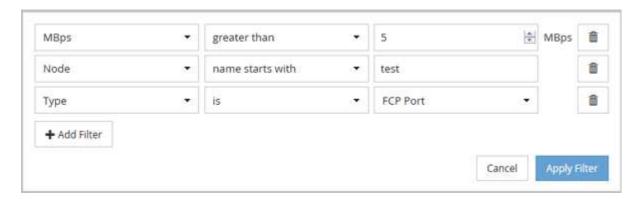
You can add filters using the Filtering panel.

1. At the top of the page, click **Filtering**. The Filtering panel displays.

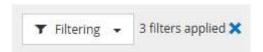
- 2. In the Filtering panel, click the left drop-down list, and select an object name: for example, *Cluster*, or a performance counter.
- 3. Click the center drop-down list, and select the boolean operator name contains or name starts with if the first selection was an object name. If the first selection was a performance counter, select greater than or less than. If the first selection was Status, select is or is not.
- 4. If your search criteria requires a numeric value, up and down arrow buttons display in the field at the right. You can click the up and down arrow buttons to display your desired numeric value.
- 5. If required, type your non-numeric search criteria in the text field at the right.
- To add filters, click Add Filter. An additional filter field displays. Complete this filter using the process
 described in the preceding steps. Note that upon adding your fourth filter, the Add Filter button no longer
 displays.
- 7. Click **Apply Filter**. The filter options are applied to the grid and an asterisk (*) is displayed in the Filtering button.
- 8. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
- 9. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example

The illustration shows the Filtering panel with three filters. The **Add Filter** button displays when you have fewer than the maximum of four filters.



After clicking **Apply Filter**, the Filtering panel closes and applies your filters.



Understanding the Unified Manager recommendations to tier data to the cloud

The Performance/Volumes inventory page displays information related to the size of the user data stored on the volume that is inactive (cold). In some cases, Unified Manager identifies certain volumes that would benefit by tiering the inactive data to the cloud tier (cloud provider or StorageGRID) of a FabricPool-enabled aggregate.



FabricPool was introduced in ONTAP 9.2, so if you are using a version of ONTAP software prior to 9.2, the Unified Manager recommendation to tier data requires upgrading your ONTAP software. Additionally, the auto tiering policy was introduced in ONTAP 9.4, so if the recommendation is to use the auto tiering policy, you must upgrade to ONTAP 9.4 or greater.

The following three fields on Performance/Volumes inventory page provide information about whether you can improve your storage system's disk utilization and save space on the performance tier by moving inactive data to the cloud tier.

Tiering Policy

The tiering policy determines whether the data on the volume remains on the performance tier or whether some of the data is moved from the performance tier to the cloud tier.

The value in this field indicates the tiering policy set on the volume, even if the volume does not currently reside on a FabricPool aggregate. The tiering policy takes effect only when the volume is on a FabricPool aggregate.

· Cold Data

The cold data displays the size of the user data stored on the volume that is inactive (cold).

A value is displayed here only when using ONTAP 9.4 or greater software because it requires that the aggregate on which the volume is deployed has the inactive data reporting parameter set to enabled, and that the minimum number of cooling days threshold has been met (for volumes that use the snapshot-only or auto tiering policy). Otherwise the value is listed as "N/A".

Cloud Recommendation

After enough information has been captured about the data activity on the volume, Unified Manager may determine there is no action required, or that you could save space on the performance tier by tiering inactive data to the cloud tier.



The Cold Data field is updated every 15 minutes, but the Cloud Recommendation field is updated every 7 days when the cold data analysis is performed on the volume. Therefore, the exact amount of cold data may differ between the fields. The Cloud Recommendation field displays the date when the analysis was run.

When Inactive Data Reporting is enabled, the Cold Data field displays the exact amount of inactive data. Without the inactive data reporting capability Unified Manager uses performance statistics to determine if data is inactive on a volume. The amount of inactive data is not displayed in the Cold Data field in this case, but it is displayed when you hover your cursor over the word **Tier** to view the cloud recommendation.

The cloud recommendations you will see are:

- Learning. Not enough data has been collected to make a recommendation.
- Tier. Analysis has determined that the volume contains inactive (cold) data and that you should configure
 the volume to move that data to the cloud tier. In some cases this may require that you move the volume to
 a FabricPool-enabled aggregate first. In other cases where the volume is already on a FabricPool
 aggregate, you just have to change the tiering policy.
- **No Action**. Either the volume has very little inactive data, the volume is already set to the "auto" tiering policy on a FabricPool aggregate, or the volume is a data protection volume. This value is also displayed

when the volume is offline or when it is being used in a MetroCluster configuration.

To move a volume, or to change the volume tiering policy or the aggregate inactive data reporting settings, use OnCommand System Manager, the ONTAP CLI commands, or a combination of these tools.

If you are logged in to Unified Manager with the OnCommand Administrator or Storage Administrator role, the **Configure Volume** link is available in the cloud recommendation when you hover your cursor over the word **Tier**. Click this button to open the Volumes page in System Manager to make the recommended change.

Descriptions of the Performance inventory pages

You use the Performance inventory pages to see a summary of performance information about each of the available storage objects, such as clusters, aggregates, volumes, and so on. You can link to the Performance object detail pages to view detailed information for a particular object.

Performance/Clusters inventory page

The Performance/Clusters inventory page displays an overview of the performance events, data, and configuration information for each cluster that is monitored by an instance of Unified Manager. This page enables you to monitor the performance of your clusters, and to troubleshoot performance issues and threshold events.

Depending on how you navigate to this page, a different title may be displayed on the page to indicate whether the list has been filtered. For example, when displaying all clusters, the title is "Clusters". When displaying a subset of clusters that are returned from the Threshold Policies page, the title is "Clusters on which policy XYZ is applied".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.

Clusters inventory page columns

The Performance/Clusters inventory page contains the following columns for each cluster.

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

Cluster

The name of the cluster. You can click the cluster name to navigate to that cluster's performance details page.

IOPS

The input/output operations per second on the cluster.

MBps

The throughput on the cluster, measured in megabytes per second.

Free Capacity

The unused storage capacity for this cluster, in gigabytes.

Total Capacity

The total storage capacity for this cluster, in gigabytes.

Node Count

The number of nodes in the cluster. You can click the node count number to navigate to the Performance/Nodes inventory page.

Host Name or IP Address

The host name or IP address (IPv4 or IPv6) of the cluster management LIF.

· Serial

The unique identification number of the cluster.

OS Version

The version of ONTAP software that is installed on the cluster.



If different versions of ONTAP software are installed on the nodes in the cluster, the lowest version number is listed. You can view the ONTAP version that is installed on each node from the Performance/Nodes inventory page.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Performance/Nodes inventory page

The Performance/Nodes inventory page displays an overview of the performance events, data, and configuration information for each node that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your nodes, and to troubleshoot performance issues and threshold events.

Depending on how you navigate to this page, the top of the page may display a different title to indicate whether the list has been filtered. For example, when displaying all nodes, the title is "Nodes". When displaying a subset of nodes that is returned from the Cluster Inventory page, the title is "Nodes on cluster: opm-cluster2".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.

Nodes inventory page columns

The Performance/Nodes inventory page contains the following columns for each node:

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

Node

The name of the node. You can click the node name to navigate to that node's performance details page.

Latency

The average response time for all I/O requests on the node, expressed in milliseconds per operation.

· IOPS

The average input/output operations per second on the node.

MBps

The throughput on the node, measured in megabytes per second.

• Flash Cache Reads

The percentage of read operations on the node that are satisfied by cache, instead of being returned from the disk.



Flash Cache data is displayed only for nodes, and only when a Flash Cache module is installed in the node.

Performance Capacity Used

The percentage of performance capacity that is being consumed by the node.



Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

Utilization

Indicates whether the CPU or memory on the node is being overused.

Free Capacity

The unused storage capacity of the node, in gigabytes.

Total Capacity

The total storage capacity of the node, in gigabytes.

Cluster

The cluster to which the node belongs. You can click the cluster's name to navigate to that cluster's details page.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Performance/Aggregates inventory page

The Performance/Aggregates inventory page displays an overview of the performance events, data, and configuration information for each aggregate that is monitored by an instance of Unified Manager. This page enables you to monitor the performance of your aggregates, and to troubleshoot performance issues and threshold events.

Depending on how you navigate to this page, a different title may be displayed on the page to indicate whether the list has been filtered. For example, when displaying all aggregates, the title is "Aggregates". When displaying a subset of aggregates that are returned from the Threshold Policies page, the title is "Aggregates on which policy aggr_IOPS is applied".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.

Root aggregates are not displayed on this page.

Aggregates inventory page columns

The Performance/Aggregates inventory page contains the following columns for each aggregate.

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

Aggregate

You can click the aggregate name to navigate to that aggregate's performance details page.

Aggregate Type

The type of aggregate:

- HDD
- · Hybrid

Combines HDDs and SSDs, but Flash Pool has not been enabled.

Hybrid (Flash Pool)

Combines HDDs and SSDs, and Flash Pool has been enabled.

- · SSD
- SSD (FabricPool)

Combines SSDs and a cloud tier

VMDisk (SDS)

Virtual disks within a virtual machine

VMDisk (FabricPool)

Combines virtual disks and a cloud tier

 LUN (FlexArray)
 This column displays "Not Available" when the monitored storage system is running a version of ONTAP earlier than 8.3.

Latency

The average response time for all I/O requests on the aggregate, expressed in milliseconds per operation.

· IOPS

The input/output operations per second on the aggregate.

MBps

The throughput on the aggregate, measured in megabytes per second.

Performance Capacity Used

The percentage of performance capacity that is being used by the aggregate.



Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

Utilization

The percentage of the aggregate's disks that are currently being used.

Free Capacity

The unused storage capacity for this aggregate, in gigabytes.

Total Capacity

The total storage capacity for this aggregate, in gigabytes.

Inactive Data Reporting

Whether the inactive data reporting capability is enabled or disabled on this aggregate. When enabled, volumes on this aggregate display the amount of cold data in the Performance/Volumes inventory page.

The value in this field is "N/A" when the version of ONTAP does not support inactive data reporting.

Cluster

The cluster to which the aggregate belongs. You can click the cluster name to navigate to that cluster's details page.

Node

The node to which the aggregate belongs. You can click the node name to navigate to that node's details page.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Performance/Volumes inventory page

The Performance/Volumes inventory page displays an overview of the performance events, counter data, and configuration information for each FlexVol volume and FlexGroup volume that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your volumes, and to troubleshoot performance issues and threshold events.

Depending on how you navigate to this page, the top of the page may display a different title to indicate whether the list has been filtered. For example, when displaying all volumes, the title is "Volumes". When displaying a subset of volumes that is returned from the Threshold Policies page, the title is "Volumes on which policy: vol_IOPS is applied".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.



For data protection (DP) volumes, only counter values for user-generated traffic are displayed. When the cluster is installed with a version of ONTAP software prior to 8.3, no counter values are displayed.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.



Root volumes are not displayed on this page.

Volume inventory page columns

The Performance/Volumes inventory page contains the following columns for each volume:

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

Volume

The volume name. You can click the volume name to navigate to the volume's performance details page.

Style

The style of volume; either FlexVol or FlexGroup.

Latency

For FlexVol volumes, this is the average response time of the volume for all I/O requests, expressed in milliseconds per operation. For FlexGroup volumes, this is the average latency of all constituent volumes.

· IOPS

For FlexVol volumes, this is the number of input/output operations per second for the volume. For FlexGroup volumes, this is the sum of IOPS for all constituent volumes.

MBps

For FlexVol volumes, this is the throughput on the volume, measured in megabytes per second. For FlexGroup volumes, this is the sum of MBps for all constituent volumes.

Free Capacity

The unused storage capacity of the volume, expressed in gigabytes.

Total Capacity

The total storage capacity of the volume, expressed in gigabytes.

Tiering Policy

The tiering policy set on the volume. The policy takes affect only when the volume is deployed on a FabricPool aggregate. The available policies are:

- None. The data for this volume always remains on the performance tier.
- Snapshot Only. Only Snapshot data is moved automatically to the cloud tier. All other data remains on the performance tier.
- Backup. On data protection volumes, all transferred user data starts in the cloud tier, but later client reads can cause hot data to move back to the performance tier.
- Auto. Data on this volume is moved between the performance tier and the cloud tier automatically when ONTAP determines that the data is "hot" or "cold".

Cold Data

The size of the user data stored on the volume that is inactive (cold).

The value is listed as "N/A" in the following situations:

- When "inactive data reporting" is disabled on the aggregate on which the volume resides.
- When "inactive data reporting" is enabled, but the minimum number of days for collecting data has not been reached.
- When using the "backup" tiering policy, or when using a version of ONTAP prior to 9.4 (when inactive data reporting is not available).

Cloud Recommendation

Unified Manager runs capacity analysis on each volume to determine if you can improve your storage system's disk utilization and save space on the performance tier by moving inactive (cold) data to the cloud tier. When the recommendation is "Tier", hover your cursor over the word **Tier** to view the recommendation. Possible recommendations are:

- Learning. Not enough data has been collected to make a recommendation.
- Tier. Analysis has determined that the volume contains inactive (cold) data and that you should configure the volume to move that data to the cloud tier.
- No Action. Either the volume has very little inactive data, or the volume is already set to the "auto" tiering policy, or the version of ONTAP does not support FabricPool.
 If you are logged in to Unified Manager with the OnCommand Administrator or Storage Administrator role, when you hover your cursor over the word Tier the Configure Volume link is available to launch System Manager so you can make the recommended change.

Cluster

The cluster to which the volume belongs. You can click the cluster name to navigate to that cluster's details page.

Node

The name of the node on which the FlexVol volume resides, or the number of nodes on which the FlexGroup volume resides.

For FlexVol volumes, you can click the name to display node details in the Node details page. For FlexGroup volumes, you can click the number to display the nodes that are used in the FlexGroup in the Nodes inventory page.

SVM

The storage virtual machine (SVM) to which the volume belongs. You can click the SVM name to navigate to that SVM's details page.

Aggregate

The name of the aggregate on which the FlexVol volume resides, or the number of aggregates on which the FlexGroup volume resides.

For FlexVol volumes, you can click the name to display aggregate details in the Aggregate details page. For FlexGroup volumes, you can click the number to display the aggregates that are used in the FlexGroup in the Aggregates inventory page.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Performance/Ports inventory page

The Performance/Ports inventory page displays an overview of the performance events, data, and configuration information for each port that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your ports, and to troubleshoot performance issues and threshold events.



Performance counter values are displayed for physical ports only. Counter values are not displayed for VLANs or interface groups.

Depending on how you navigate to this page, the top of the page may display a different title to indicate whether the list has been filtered. For example, when displaying all ports, the title is "Ports". When displaying a subset of ports that is returned from the Threshold Policies page, the title is "Ports on which policy: port_IOPS is applied".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.

Ports inventory page columns

The Performance/Ports inventory page contains the following columns for each port:

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

Port

You can click the port name to navigate to that port's performance details page.

Type

The port type is either Network or Fibre Channel Protocol (FCP).

MBps

The throughput on the port, measured in megabytes per second.

Utilization

The percentage of the port's available bandwidth that is currently being used.

Cluster

The cluster to which the port belongs. You can click the cluster name to navigate to that cluster's details page.

Node

The node to which the port belongs. You can click the node name to navigate to that node's details page.

Speed

The maximum data transfer rate for the port.

Role

The network port function: either Data or Cluster. FCP ports cannot have a role, and the role is displayed as N/A.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Performance/SVMs inventory page

The Performance/SVMs inventory page displays an overview of the performance events, data, and configuration information for each storage virtual machine (SVM) that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your SVMs, and to troubleshoot performance issues and threshold events.

Depending on how you navigate to this page, the top of the page may display a different title to indicate whether the list has been filtered. For example, when displaying all SVMs, the title is "SVMs". When displaying a subset of SVMs that is returned from the Threshold Policies page, the title is "SVMs on which policy: SVM_IOPS is applied".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.



The SVMs that are listed on this page include only Data and Cluster SVMs. Unified Manager does not use or display Admin or Node SVMs.

SVMs inventory page columns

The Performance/SVMs inventory page contains the following columns for each SVM:

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

• SVM

You can click the SVM name to navigate to that SVM's performance details page.

Latency

The average response time for all I/O requests, expressed in milliseconds per operation.

· IOPS

The input/output operations per second for the SVM.

MBps

The throughput on the SVM, measured in megabytes per second.

Free Capacity

The unused storage capacity of the SVM, in gigabytes.

Total Capacity

The total storage capacity of the SVM, in gigabytes.

Cluster

The cluster to which the SVM belongs. You can click the cluster name to navigate to that cluster's details page.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Performance/LUNs inventory page

The Performance/LUNs inventory page displays an overview of the performance events, data, and configuration information for each LUN that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your LUNs, and to troubleshoot performance issues and threshold events.

Depending on how you navigate to this page, the top of the page may display a different title to indicate whether the list has been filtered. For example, when displaying all LUNs, the title is "LUNs". When displaying a subset of LUNs that is returned from the Threshold Policies page, the title is "LUNs on which policy: LUN_IOPS is applied".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.

LUNs inventory page columns

The Performance/LUNs inventory page contains the following columns for each LUN:

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

• LUN

You can click the LUN name to navigate to that LUN's performance details page.

Latency

The average response time for all I/O requests, expressed in milliseconds per operation.

· IOPS

The input/output operations per second for the LUN.

MBps

The throughput on the LUN, measured in megabytes per second.

Free Capacity

The unused storage capacity of the LUN, in gigabytes.

Total Capacity

The total storage capacity of the LUN, in gigabytes.

Cluster

The cluster to which the LUN belongs. You can click the cluster name to navigate to that cluster's details page.

Node

The node to which the LUN belongs. You can click the node name to navigate to that node's details page.

· SVM

The storage virtual machine (SVM) to which the LUN belongs. You can click the SVM name to navigate to that SVM's details page.

Aggregate

The aggregate to which the LUN belongs. You can click the aggregate name to navigate to that aggregate's details page.

Volume

The volume to which the LUN belongs. You can click the volume name to navigate to that volume's details

page.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Namespaces inventory page

The Namespaces inventory page displays an overview of the performance events, data, and configuration information for each Namespace that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance and health of your Namespaces, and to troubleshoot issues and threshold events.

Depending on how you navigate to this page, the top of the page may display a different title to indicate whether the list has been filtered. For example, when displaying all Namespaces, the title is "Namespaces". When displaying a subset of Namespaces that is returned from the Threshold Policies page, the title is "Namespaces on which policy: Namespace IOPS is applied".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.

Namespaces inventory page columns

The Namespaces inventory page contains the following columns for each Namespace:

Subsystem

The subsystem of the Namespace.

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

Namespace

You can click the Namespace name to navigate to that Namespace's performance details page.

State

The current state of the Namespace.

- Offline Read or write access to the Namespace is not allowed.
- Online Read and write access to the Namespace is allowed.
- NVFail The Namespace was automatically taken offline due to an NVRAM failure.
- Space Error The Namespace has run out of space.

SVM

The storage virtual machine (SVM) to which the Namespace belongs. You can click the SVM name to navigate to that SVM's details page.

Cluster

The cluster to which the Namespace belongs. You can click the cluster name to navigate to that cluster's details page.

Volume

The volume to which the Namespace belongs. You can click the volume name to navigate to that volume's details page.

Total Capacity

The total storage capacity of the Namespace, in gigabytes.

Free Capacity

The unused storage capacity of the Namespace, in gigabytes.

IOPS

The input/output operations per second for the Namespace.

Latency

The average response time for all I/O requests on the Namespace, expressed in milliseconds per operation.

MBps

The throughput on the Namespace, measured in megabytes per second.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Performance/LIFs inventory page

The Performance/LIFs inventory page displays an overview of the performance events, data, and configuration information for each LIF that is being monitored by this instance of Unified Manager. This page enables you to quickly monitor the performance of your LIFs, and to troubleshoot performance issues and threshold events.

Depending on how you navigate to the Performance/LIFs inventory page, the top of the page may display a different title to indicate whether the list has been filtered. For example, when displaying all LIFs, the title is "LIFs". When displaying a subset of LIFs that is returned from the Threshold Policies page, the title is "LIFs on which policy: LIF IOPS is applied".

The buttons along the top of the page enable you to perform searches to locate specific data, create and apply filters to narrow the list of displayed data, export the data on the page to a .csv file, and add or remove columns from the page.

By default, objects on the object inventory pages are sorted based on object performance event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. The values of the performance counters are based on an average from the previous 72 or more hours of data, as indicated on the page. You can click the refresh button to update the object inventory data.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear PerformanceThreshold Policy** buttons.



The LIFs that are listed on the Performance/LIFs inventory page include Data LIFs, Cluster LIFs, Node Management LIFs, and Intercluster LIFs. Unified Manager does not use or display System LIFs.

LIFs inventory page columns

The Performance/LIFs inventory page contains the following columns for each LIF.

Status

A healthy object with no active events displays a green check mark icon (\bigcirc). If the object has an active event, the event indicator icon identifies the event severity: critical events are red (\bigcirc), error events are orange (\bigcirc), and warning events are yellow (\bigcirc).

• LIF

You can click the LIF name to navigate to the performance details page of that LIF.

Type

The LIF type: Network (iSCSI, NFS, CIFS), FCP, or NVMf FCP.

Latency

The average response time for all I/O requests, expressed in milliseconds per operation. Latency is not applicable to NFS LIFs and CIFS LIFs, and is displayed as N/A for these types.

· IOPS

The input/output operations per second. IOPS is not applicable to NFS LIFs and CIFS LIFs, and is displayed as N/A for these types.

MBps

The throughput on the LIF, measured in megabytes per second.

Cluster

The cluster to which the LIF belongs. You can click the cluster's name to navigate to that cluster's details page.

· SVM

The storage virtual machine to which the LIF belongs. You can click the SVM name to navigate to that SVM's details page.

Home Location

The home location for the LIF, displayed as node name and port name, separated by a colon (:). If the location is displayed with an ellipsis (...), you can position your cursor over the location name to view the full location.

Current Location

The current location for the LIF, displayed as node name and port name, separated by a colon (:). If the location is displayed with an ellipsis (...), you can position your cursor over the location name to view the full location.

Role

The LIF role: Data, Cluster, Node Management, or Intercluster.

Threshold Policy

The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (...) to view the full policy name or the list of assigned policy names. The **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Monitoring performance using the Performance Explorer pages

The Performance Explorer pages display detailed information about the performance of each object in a cluster. The page provides a detailed view into the performance of all cluster objects, enabling you to select and compare the performance data of specific objects across various time periods.

You can also assess the overall performance of all objects, and compare object performance data in a side-by-side format.

If an object is no longer managed by Unified Manager, the status Removed is displayed to the right of the

object's name at the top of the Performance Explorer page.

Understanding the root object

The root object is the baseline against which other object comparisons are made. This enables you to view and compare the data from other objects to the root object, providing performance data analysis that helps you to troubleshoot and improve object performance.

The root object name displays at the top of the Comparing pane. Additional objects display below the root object. Although there is no limit to the number of additional objects you can add to the Comparing pane, only one root object is allowed. Data for the root object automatically displays in the graphs in the Counter Charts pane.

You cannot change the root object; it is always set to the object page you are viewing. For example, if you open the Volume Performance Explorer page of Volume1, then Volume1 is the root object and cannot be changed. If you want to compare against a different root object, then you must click the link for an object and open its landing page.



Events and Thresholds are displayed only for root objects.

Apply filtering to reduce the list of correlated objects in the grid

Filtering enables you to display a smaller, more well-defined subset of objects in the grid. For example, if you have 25 volumes in the grid, filtering enables you to view only those volumes that have throughput less than 90 MBps, or latency greater than 1 ms/op.

Specifying a time range for correlated objects

The Time Range selector on the Performance Explorer page enables you to specify the time range for object data comparison. Specifying a time range refines the contents of the Performance Explorer pages to show only the object data within the time range you have specified.

About this task

Refining the time range provides an efficient method of displaying only the performance data in which you are interested. You can select a predefined time range or specify a custom time range. The default time range is the preceding 72 hours.

Selecting a predefined time range

Selecting a predefined time range is a quick and efficient way for you to customize and focus data output when viewing cluster object performance data. When selecting a predefined time range, data for up to 13 months is available.

Steps

1. At the top right of the **Performance Explorer** page, click **Time Range**.

- From the right side of the Time Range Selection panel, select a predefined time range.
- 3. Click Apply Range.

Specifying a custom time range

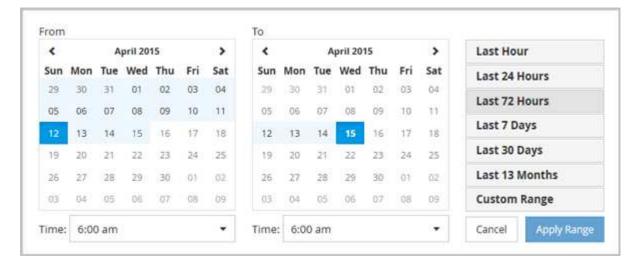
The Performance Explorer page enables you to specify the date and time range for your performance data. Specifying a custom time range provides greater flexibility than using predefined time ranges when refining cluster object data.

About this task

You can select a time range between one hour and 390 days. 13 months equals 390 days because each month is counted as 30 days. Specifying a date and time range provides more detail and enables you to zoom in on specific performance events or series of events. Specifying a time range also aids in troubleshooting potential performance issues, as specifying a date and time range displays data surrounding the performance event in finer detail. Use the **Time Range** control to select predefined date and time ranges, or specify your own custom date and time range of up to 390 days. Buttons for predefined time ranges vary from the **Last Hour** through the **Last 13 Months**.

Selecting the **Last 13 Months** option or specifying a custom date range greater than 30 days displays a dialog box alerting you that performance data displayed for a period greater than 30 days is charted using hourly averages and not 5-minute data polling. Therefore, a loss of timeline visual granularity might occur. If you click the **Do not show again** option in the dialog box, the message does not appear when you select the **Last 13 Months** option or specify a custom date range greater than 30 days. Summary data also applies on a smaller time range, if the time range includes a time/date that is more than 30 days from today.

When selecting a time range (either custom or predefined), time ranges of 30 days or fewer are based on 5-minute interval data samples. Time ranges greater than 30 days are based on one-hour interval data samples.



- 1. Click the **Time Range** drop-down box and the Time Range panel displays.
- 2. To select a predefined time range, click one of the **Last...** buttons at the right of the **Time Range** panel. When selecting a predefined time range, data for up to 13 months is available. The predefined time range button you selected is highlighted, and the corresponding days and time display in the calendars and time selectors.
- 3. To select a custom date range, click the start date in the **From** calendar on the left. Click < or > to navigate forward or backward in the calendar. To specify the end date, click a date in the **To** calendar on the right. Note that the default end date is today unless you specify a different end date. The **Custom Range** button

at the right of the Time Range panel is highlighted, indicating that you have selected a custom date range.

- 4. To select a custom time range, click the **Time** control below the **From** calendar and select the start time. To specify the end time, click the **Time** control below the **To** calendar on the right and select the end time. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom time range.
- 5. Optionally, you can specify the start and end times when selecting a predefined date range. Select the predefined date range as previously described, then select the start and end times as previously described. The selected dates are highlighted in the calendars, your specified start and end times display in the **Time** controls, and the **Custom Range** button is highlighted.
- 6. After selecting the date and time range, click **Apply Range**. The performance statistics for that time range display in the charts and in the Events timeline .

Defining the list of correlated objects for comparison graphing

You can define a list of correlated objects for data and performance comparison in the Counter Chart pane. For example, if your storage virtual machine (SVM) is experiencing a performance issue, you can compare all volumes in the SVM to identify which volume might be causing the issue.

About this task

You can add any object in the correlated objects grid to the Comparing and Counter Chart panes. This enables you to view and compare data of multiple objects and with the root object. You can add and remove objects to and from the correlated objects grid; however, the root object in the Comparing pane is not removable.



Adding many objects to the Comparing pane may have a negative impact on performance. To maintain performance, you should select a limited number of charts for data comparison.

Steps

1. In the objects grid, locate the object that you want to add, and click the **Add** button.

2. Hide or show data for selected objects:

To do this	Take this action
Hide a selected object	Click the selected object's eye icon () in the Comparing pane. The object's data is hidden, and the eye icon for that object turns gray.
Show a hidden object	Click the gray eye icon of the selected object in the Comparing pane. The eye icon returns to its original color, and the object data is added back into the graphs in the Counter Charts pane.

3. Remove selected objects from the **Comparing** pane:

To do this	Take this action
Remove a selected object	Hover over the selected object's name in the Comparing pane to show the remove object button (X), and then click the button. The object is removed from the Comparing pane, and its data is cleared from the counter charts.
Remove all selected objects	Click the remove all object's button (X) at the top of the Comparing pane. All selected objects and their data are removed, leaving only the root object.

Understanding counter charts

Charts in the Counter Charts pane enable you to view and compare performance data for the root object and for objects you have added from the correlated objects grid. This can help you understand performance trends and isolate and resolve performance issues.

Counter charts displayed by default are Events, Latency,IOPS, and MBps. Optional charts that you can choose to display are Utilization, Performance Capacity Used, Available IOPS, IOPS/TB, and Cache Miss Ratio. Additionally, you can choose to view total values or breakdown values for the Latency, IOPS, MBps, and Performance Capacity Used charts.

The Performance Explorer displays certain counter charts by default; whether the storage object supports them all or not. When a counter is not supported, the counter chart is empty and the message Not applicable for <object> is displayed.

The charts display performance trends for the root object and for all objects you have selected in the Comparing pane. Data in each chart is arranged as follows:

X axis

Displays the specified time period. If you have not specified a time range, the default is the preceding 72-hour period.

Y axis

Displays counter units unique to the selected object, or objects.

Trend line colors match the color of the object name as displayed in the Comparing pane. You can position your cursor over a point on any trend line to view details for time and value for that point.

If you want to investigate a specific period of time within a chart, you can use one of the following methods:

- Use the < button to expand the Counter Charts pane to span the width of the page.
- Use the cursor (when it transitions to a magnifying glass) to select a portion of the timeframe in the chart to focus and enlarge that area. You can click Reset Chart Zoom to return the chart to the default timeframe.
- Use the **Zoom View** button to display a large single counter chart that contains expanded details and threshold indicators.



Occasionally, gaps in the trend lines display. Gaps mean that either Unified Manager failed to collect performance data from the storage system or that Unified Manager might have been down.

Types of performance counter charts

There are standard performance charts that display the counter values for the selected storage object. Each of the Breakdown counter charts display the total values separated out into read, write, and other categories. Furthermore, some Breakdown counter charts display additional detail when the chart is displayed in Zoom view.

The following table shows the available performance counter charts.

Available charts	Chart description
Events	Displays critical, error, warning, and information events in correlation with the statistical charts for the root object. Health events display in addition to performance events to provide a complete picture of the reasons performance may be affected.
Latency - Total	Number of milliseconds required to respond to application requests. Note that the average latency values are I/O weighted.
Latency - Breakdown	The same information shown in Latency Total, but with the performance data separated into read, write, and other latency. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace.
Latency - Cluster Components	The same information shown in Latency Total, but with the performance data separated into latency by cluster component. This chart option applies only when the selected object is a volume.
IOPS - Total	Number of input/output operations processed per second.
IOPS - Breakdown	The same information shown in IOPS Total, but with the performance data separated into read, write, and other IOPS. When displayed in Zoom view the volumes chart displays QoS minimum and maximum throughput values, if configured in ONTAP. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace.

Available charts	Chart description		
IOPS - Protocols	The same information shown in IOPS Total, but the performance data is separated into individual charts for CIFS, NFS, FCP, NVMe, and iSCSI protocol traffication chart option applies only when the selected object is an SVM.		
IOPS/TB - Total	Number of input/output operations processed per second based on the total space that is being consumed by the workload, in terabytes. Also called I/O density, this counter measures how much performance can be delivered by a given amount of storage capacity. When displayed in Zoom view the volumes chart displays QoS expected and peak throughput values, if configured in ONTAP. This chart option applies only when the selected object is a volume.		
MBps - Total	Number of megabytes of data transferred to and from the object per second.		
MBps - Breakdown	The same information shown in the MBps chart, but with the MBps data separated into disk reads, Flash Cache reads, writes, and other. When displayed in Zoom view, the volumes chart displays QoS maximum throughput values, if configured in ONTAP. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace. Flash Cache data is displayed only for nodes, and only when a Flash Cache module is installed in the node.		
Performance Capacity Used - Total	Percentage of performance capacity that is being consumed by the node or aggregate. Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.		
Performance Capacity Used - Breakdown	Performance Capacity Used data separated into user protocols and system background processes. Additionally, the amount of free performance capacity is shown.		

Available charts	Chart description	
Available IOPS - Total	Number of input/output operations per second that are currently available (free) on this object. This number is the result of subtracting the currently used IOPS from the total IOPS that Unified Manager calculates that the object can perform. This chart option applies only when the selected object is a nod or aggregate. Available IOPS data is available only	
	when the nodes in a cluster are installed with ONTAP 9.0 or later software.	
Utilization - Total	Available resource percentage of the object that is being used. Utilization indicates node utilization for nodes, disk utilization for aggregates, and bandwidth utilization for ports. This chart option applies only when the selected object is a node, aggregate, or port.	
Cache Miss Ratio - Total	Percentage of read requests from client applications that are returned from the disk instead of being returned from the cache. This chart option applies only when the selected object is a volume.	

Selecting performance charts to display

The Choose charts drop-down list enables you to select the types of performance counter charts to display in the Counter Charts pane. This enables you to view specific data and counters, based on your performance requirements.

Steps

- 1. In the Counter Charts pane, click the Choose charts drop-down list.
- 2. Add or remove charts:

То	Do this
Add or remove individual charts	Click the check boxes next to the charts you want to show or hide
Add all charts	Click Select All
Remove all charts	Click Unselect All

Your chart selections are displayed in the Counter Charts pane. Note that as you add charts, the new charts are inserted into the Counter Charts pane to match the order of the charts listed in the Choose

Expanding the Counter Charts pane

You can expand the Counter Charts pane so that the charts are larger and more readable.

About this task

After you have defined the comparison objects and the time range for counters, you can view a larger Counter Charts pane. You use the < button in the middle of the Performance Explorer window to expand the pane.

Steps

1. Expand or reduce the Counter Charts pane.

То	Do this
Expand the Counter Charts pane to fit the width of the page	Click the < button
Reduce the Counter Charts pane to the right half of the page	Click the > button

Changing the Counter Charts focus to a shorter period of time

You can use your mouse to reduce the time range to focus on a specific period of time in the Counter Chart pane or in the Counter Charts Zoom View window. This enables you to see a more granular and microscopic view of any part of the timeline of performance data, events, and thresholds.

Before you begin

The cursor must have changed to a magnifying glass to indicate that this functionality is active.



When using this feature, which alters the timeline to display values that correspond to the more granular display, the time and date range on the **Time Range** selector does not change from the original values for the chart.

Steps

1. To zoom into a specific period of time, click using the magnifying glass and drag the mouse to highlight the area that you want to see in detail.

The counter values for the time period you select fills the counter chart.

To return to the original period of time as set in the Time Range selector, click the Reset Chart Zoom button.

The counter chart displays in its original state.

Viewing event details in the Events Timeline

You can view all events and their related details in the Events Timeline pane of Performance Explorer. This is a quick and efficient method of viewing all the health and performance events that occurred on the root object during a specified time range, which can be helpful when troubleshooting performance issues.

About this task

The Events Timeline pane shows critical, error, warning, and informational events that occurred on the root object during the selected time range. Each event severity has its own timeline. Single and multiple events are represented by an event dot on the timeline. You can position your cursor over an event dot to see the event details. To increase the visual granularity of multiple events, you can decrease the time range. This spreads out multiple events into single events, enabling you to separately view and investigate each event.

Each performance event dot on the Events Timeline lines up vertically with a corresponding spike in the counter charts trend lines that are displayed below the Events Timeline. This provides a direct visual correlation between events and overall performance. Health events are displayed on the timeline as well, but these types of events do not necessarily line up with a spike in one of the performance charts.

Steps

1. On the **Events Timeline** pane, position the cursor over an event dot on a timeline to view a summary of the event or events at that event point.

A pop-up dialog displays information about the event types, the date and time when the events occurred, the state, and the event duration.

2. View full event details for one event or multiple events:

To do this	Click this	
View details for a single event	View Event Detail in the pop-up dialog.	
View details for multiple events	View Event Details in the pop-up dialog.	
	i	Clicking a single event on the multiple events dialog displays the appropriate Event Details page.

Counter Charts Zoom View

The Counter Charts provide a Zoom View that enables you to zoom in on performance details over your specified time period. This enables you to see performance details and events with much higher granularity, which is beneficial when troubleshooting performance issues.

When displayed in Zoom View, some of the breakdown charts provide additional information than what appears when the chart is not in Zoom View. For example, the IOPS, IOPS/TB, and MBps Breakdown chart Zoom View pages display QoS policy values for volumes and LUNs if they have been set in ONTAP.



For system-defined performance threshold policies, only the "Node resources over-utilized" and "QoS throughput limit breached" policies are available from the **Policies** list. The other system-defined threshold policies are not available at this time.

Displaying the Counter Charts Zoom View

The Counter Charts Zoom View provides a finer level of detail for the selected counter chart and its associated timeline. This magnifies the counter chart data, enabling you to have a sharper view into performance events and their underlying causes.

About this task

You can display the Counter Charts Zoom View for any counter chart.

Steps

- 1. Click **Zoom View** to open the selected chart a new browser window.
- 2. If you are viewing a Breakdown chart and then click **Zoom View** the Breakdown chart is shown in Zoom View. You can select **Total** while in Zoom View if you want to change the view option.

Specifying the time range in Zoom View

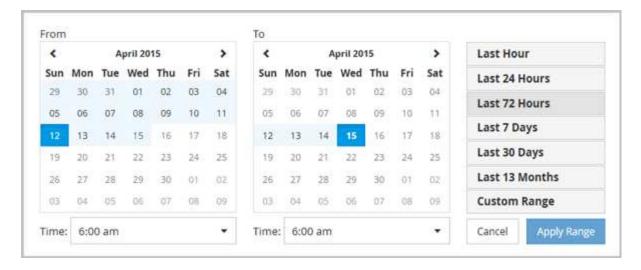
The **Time Range** control in the Counter Charts Zoom View window enables you to specify a date and time range for the selected chart. This enables you to quickly locate specific data based on either a preset time range or your own custom time range.

About this task

You can select a time range between one hour and 390 days. 13 months equals 390 days because each month is counted as 30 days. Specifying a date and time range provides more detail and enables you to zoom in on specific performance events or series of events. Specifying a time range also aids in troubleshooting potential performance issues, as specifying a date and time range displays data surrounding the performance event in finer detail. Use the **Time Range** control to select predefined date and time ranges, or specify your own custom date and time range of up to 390 days. Buttons for predefined time ranges vary from the **Last Hour** through the **Last 13 Months**.

Selecting the **Last 13 Months** option or specifying a custom date range greater than 30 days displays a dialog box alerting you that performance data displayed for a period greater than 30 days is charted using hourly averages and not 5-minute data polling. Therefore, a loss of timeline visual granularity might occur. If you click the **Do not show again** option in the dialog box, the message does not appear when you select the **Last 13 Months** option or specify a custom date range greater than 30 days. Summary data also applies on a smaller time range, if the time range includes a time/date that is more than 30 days from today.

When selecting a time range (either custom or predefined), time ranges of 30 days or fewer are based on 5-minute interval data samples. Time ranges greater than 30 days are based on one-hour interval data samples.



- 1. Click the **Time Range** drop-down box and the Time Range panel displays.
- To select a predefined time range, click one of the Last... buttons at the right of the Time Range panel.
 When selecting a predefined time range, data for up to 13 months is available. The predefined time range
 button you selected is highlighted, and the corresponding days and time display in the calendars and time
 selectors.
- 3. To select a custom date range, click the start date in the **From** calendar on the left. Click < or > to navigate forward or backward in the calendar. To specify the end date, click a date in the **To** calendar on the right. Note that the default end date is today unless you specify a different end date. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom date range.
- 4. To select a custom time range, click the **Time** control below the **From** calendar and select the start time. To specify the end time, click the **Time** control below the **To** calendar on the right and select the end time. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom time range.
- 5. Optionally, you can specify the start and end times when selecting a predefined date range. Select the predefined date range as previously described, then select the start and end times as previously described. The selected dates are highlighted in the calendars, your specified start and end times display in the **Time** controls, and the **Custom Range** button is highlighted.
- 6. After selecting the date and time range, click **Apply Range**. The performance statistics for that time range display in the charts and in the Events timeline.

Selecting performance thresholds in Counter Charts Zoom View

Applying thresholds in the Counter Charts Zoom View provides a detailed view of occurrences of performance threshold events. This enables you to apply or remove thresholds, and immediately view the results, which can be helpful while deciding whether troubleshooting should be your next step.

About this task

Selecting thresholds in the Counter Charts Zoom View enables you to view precise data about performance threshold events. You can apply any threshold that appears under the **Policies** area of the Counter Charts Zoom View.

Only one policy at a time can be applied to the object in the Counter Charts Zoom View.

Steps

1. Select or deselect the

that is associated with a policy.

The selected threshold is applied to the Counter Charts Zoom View. Critical thresholds are displayed as a red line; warning thresholds are displayed as a yellow line.

Viewing workload QoS minimum and maximum settings

You can view the ONTAP-defined quality of service (QoS) policy settings on a volume or LUN in the Performance Explorer charts. A throughput maximum setting limits the impact of competing workloads on system resources. A throughput minimum setting ensures that a critical workload meets minimum throughput targets regardless of demand by competing workloads.

About this task

QoS throughput "minimum" and "maximum" IOPS and MBps settings are displayed in the counter charts only if they have been configured in ONTAP. Throughput minimum settings are available only on systems running ONTAP 9.2 or later software, only on AFF systems, and they can be set only for IOPS at this time.

Adaptive QoS policies are available starting with ONTAP 9.3 and are expressed using IOPS/TB instead of IOPS. These policies automatically adjust the QoS policy value based on the volume size, per workload, thereby maintaining the ratio of IOPS to terabytes as the size of the volume changes. You can apply an adaptive QoS policy group to volumes only. The QoS terminology "expected" and "peak" are used for adaptive QoS policies instead of minimum and maximum.

Unified Manager generates warning events for QoS policy breaches when workload throughput has exceeded the defined QoS maximum policy setting during each performance collection period for the previous hour. Workload throughput may exceed the QoS threshold for only a short period of time during each collection period, but Unified Manager displays the "average" throughput during the collection period on the chart. For this reason you may see QoS events while the throughput for a workload might not have crossed the policy threshold shown in the chart.

Steps

1. In the **Performance Explorer** page for your selected volume or LUN, perform the following actions to view the QoS ceiling and floor settings:

If you want to	Do this
View the IOPS ceiling (the QoS max)	In the IOPS Total or Breakdown chart, click Zoom View .
View the MBps ceiling (the QoS max)	In the MBps Total or Breakdown chart, click Zoom View .
View the IOPS floor (the QoS min)	In the IOPS Total or Breakdown chart, click Zoom View .

If you want to	Do this
View the IOPS/TB ceiling (the QoS peak)	For volumes, in the IOPS/TB chart, click Zoom View .
View the IOPS/TB floor (the QoS expected)	For volumes, in the IOPS/TB chart, click Zoom View .

The dashed, horizontal line indicates the maximum or minimum throughput value set in ONTAP. You can also view when changes to the QoS values were implemented.

2. To view the specific IOPS and MBps values compared to the QoS setting, move your cursor into the chart area to see the popup window.

After you finish

If you notice that certain volumes or LUNs have very high IOPS or MBps and are stressing system resources, you can use System Manager or the ONTAP CLI to adjust the QoS settings so that these workloads do not affect the performance of other workloads.

For more information on adjusting QoS settings, see the ONTAP 9 Performance Monitoring Power Guide.

ONTAP 9 Performance Monitoring Power Guide

How different types of QoS policies are displayed in Unified Manager

You can view the ONTAP-defined quality of service (QoS) policy settings that have been applied to a volume or LUN in the Performance Explorer IOPS, IOPS/TB, and MBps charts. The information displayed in the charts is different depending on the type of QoS policy that has been applied to the workload.

A throughput "ceiling" setting defines the maximum throughput that the workload can consume, and thereby limits the impact on competing workloads for system resources. A throughput "floor" setting defines the minimum throughput that must be available to the workload so that a critical workload meets minimum throughput targets regardless of demand by competing workloads.

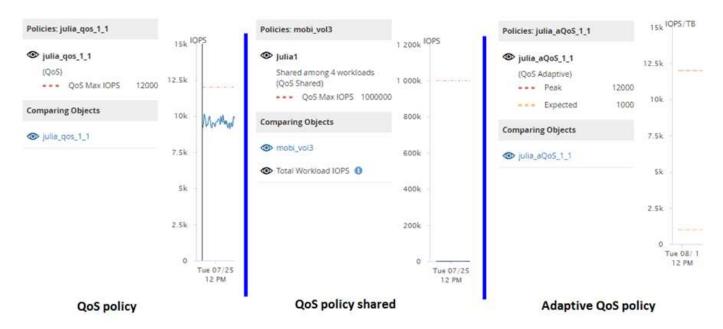
Shared and non-shared QoS policies for IOPS and MBps use the terms "minimum" and "maximum" to define the floor and ceiling. Adaptive QoS policies for IOPS/TB, which were introduced in ONTAP 9.3, use the terms "expected" and "peak" to define the floor and ceiling.

While ONTAP enables you to create these two types of QoS policies, depending on how they are applied to workloads there are three ways that the QoS policy will be displayed in the performance charts.

Type of policy	Functionality	Indicator in Unified Manager interface
QoS shared policy assigned to a single workload, or QoS non-shared policy assigned to a single workload or multiple workloads	Each workload can consume the specified throughput setting	Displays "(QoS)"

Type of policy	Functionality	Indicator in Unified Manager interface
QoS shared policy assigned to multiple workloads	All workloads share the specified throughput setting	Displays "(QoS Shared)"
Adaptive QoS policy assigned to a single workload or multiple workloads	Each workload can consume the specified throughput setting	Displays "(QoS Adaptive)"

The following figure shows an example of how the three options are shown in the counter charts.



When a normal QoS policy that has been defined in IOPS appears in the IOPS/TB chart for a workload, ONTAP converts the IOPS value to an IOPS/TB value and Unified Manager displays that policy in the IOPS/TB chart along with the text "QoS, defined in IOPS".

When an adaptive QoS policy that has been defined in IOPS/TB appears in the IOPS chart for a workload, ONTAP converts the IOPS/TB value to an IOPS value and Unified Manager displays that policy in the IOPS chart along with the text "QoS Adaptive, defined in IOPS/Used TB" or "QoS Adaptive, defined in IOPS/Allocated TB" depending on how the peak IOPS allocation setting is configured. When the allocation setting is set to "allocated-space", the peak IOPS is calculated based on the size of the volume. When the allocation setting is set to "used-space", the peak IOPS is calculated based on the amount of data stored in the volume, taking into account storage efficiencies.



The IOPS/TB chart displays performance data only when the logical capacity used by the volume is greater than or equal to 1 TB. Gaps are displayed in the chart when the used capacity falls below 1 TB during the selected timeframe.

Viewing volume latency by cluster component

You can view detailed latency information for a volume by using the Performance/Volume Explorer page. The Latency - Total counter chart shows total latency on the volume, and the Latency - Breakdown counter chart is useful for determining the impact of read and

write latency on the volume.

About this task

Additionally, the Latency - Cluster Components chart shows a detailed comparison of the latency of each cluster component to help determine how each component contributes to the total latency on the volume. The following cluster components are displayed:

- Network
- QoS Policy
- · Network Processing
- Cluster Interconnect
- · Data Processing
- · Aggregate Operations
- MetroCluster Resources
- Cloud Latency
- Sync SnapMirror

Steps

1. In the **Performance/Volume Explorer** page for your selected volume, from the Latency chart, select **Cluster Components** from the drop-down menu.

The Latency - Cluster Components chart is displayed.

To view a larger version of the chart, select Zoom View.

The cluster component comparative chart is displayed. You can restrict the comparison by deselecting or selecting the that is associated with each cluster component.

3. To view the specific values, move your cursor into the chart area to see the popup window.

Viewing SVM IOPS traffic by protocol

You can view detailed IOPS information for an SVM by using the Performance/SVM Explorer page. The IOPS - Total counter chart shows total IOPS usage on the SVM, and the IOPS - Breakdown counter chart is useful for determining the impact of read, write, and other IOPS on the SVM.

About this task

Additionally, the IOPS - Protocols chart shows a detailed comparison of the IOPS traffic for each protocol that is being used on the SVM. The following protocols are available:

- CIFS
- NFS
- FCP
- iSCSI

NVMe

Steps

1. In the **Performance/SVM Explorer** page for your selected SVM, from the IOPS chart, select **Protocols** from the drop-down menu.

The IOPS - Protocols chart is displayed.

2. To view a larger version of the chart, select **Zoom View**.

The IOPS advanced protocol comparative chart is displayed. You can restrict the comparison by deselecting or selecting the that is associated with a protocol.

3. To view the specific values, move your cursor into the chart area of either chart to see the popup window.

Viewing volume and LUN latency charts to verify performance guarantee

You can view the volumes and LUNs that you have subscribed to the "Performance Guarantee" program to verify that latency has not exceeded the level you have been guaranteed.

About this task

The latency performance guarantee is a millisecond per operation value that should not be exceeded. It is based on an hourly average, not on the default five minute performance collection period.

Steps

- In the Performance Volumes or Performance LUNs inventory page, select the volume or LUN that you
 are interested in.
- 2. In the **Performance Explorer** page for your selected volume or LUN, choose **Hourly Average** from the **View statistics in** selector.

The horizontal line in the Latency chart will show a smoother line as the five-minute collections are replaced with the hourly average.

3. If you have other volumes on the same aggregate that are under the performance guarantee, you can add those volumes to view their latency value in the same chart.

Components of the Object Landing pages

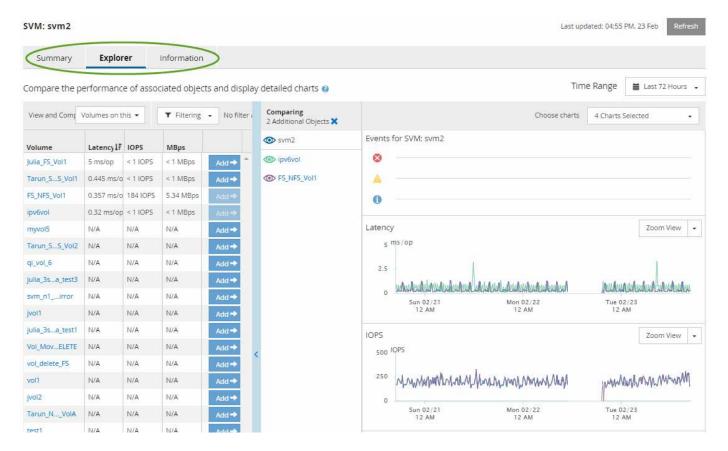
The Object Landing pages provide details about all critical, warning, and informational events. They provide a detailed view into the performance of all cluster objects, enabling you to select and compare individual objects across various time periods.

The Object Landing pages enable you to examine the overall performance of all objects, and to compare object performance data in a side-by-side format. This is beneficial when assessing performance and when troubleshooting events.



The data displayed in the counter summary panels and in the Counter Charts are based on a five-minute sampling interval. The data displayed in the objects inventory grid in the left side of the page is based on a one-hour sampling interval.

The following image shows an example of an Object Landing page displaying the Explorer information:



Depending on the storage object that is being viewed, the Object Landing page can have the following tabs that provide performance data about the object:

Summary

Displays three or four counter charts containing the events and performance per object for the preceding 72-hour period, including a trend line that shows the high and low values during that period.

Explorer

Displays a grid of storage objects that are related to the current object, which enables you to compare the performance values of the current object with those of the related objects. This tab includes up to eleven counter charts and a time range selector, which enable you to perform a variety of comparisons.

Information

Displays values for non-performance configuration attributes about the storage object, including the installed version of ONTAP software, HA partner name, and number of ports and LIFs.

Top Performers

For clusters: Displays the storage objects that have the highest performance or the lowest performance, based on the performance counter that you select.

Failover Planning

For nodes: Displays the estimate of the performance impact on a node if the HA partner of the node fails.

Details

For volumes: Displays detailed performance statistics for all I/O activity and operations for the selected volume workload. This tab is available for FlexVol volumes, FlexGroup volumes, and constituents of FlexGroups.

Summary page

The Summary page displays counter charts that contain details about the events and performance per object for the preceding 72-hour period. This data is not automatically refreshed, but is current as of the last page load. The charts in the Summary page answer the question *Do I need to look further?*

Charts and counter statistics

The summary charts provide a quick, high-level overview for the last 72-hour period, and help you to identify possible issues that require further investigation.

The Summary page counter statistics are displayed in graphs.

You can position your cursor over the trend line in a graph to view the counter values for a particular point in time. The summary charts also display the total number of active critical and warning events for the preceding 72-hour period for the following counters:

Latency

Average response time for all I/O requests; expressed in milliseconds per operation.

Displayed for all object types.

· IOPS

Average operating speed; expressed in input/output operations per second.

Displayed for all object types.

MBps

Average throughput; expressed in megabytes per second.

Displayed for all object types.

Performance Capacity Used

Percentage of performance capacity that is being consumed by a node or aggregate.

Displayed for nodes and aggregates only. This chart is displayed only when using ONTAP 9.0 or later software.

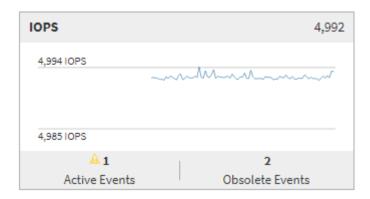
Utilization

Percentage of object utilization for nodes and aggregates, or bandwidth utilization for ports.

Displayed for nodes, aggregates, and ports only.

Positioning the cursor over the event count for Active events shows the type and number of events. Critical events are displayed in red (), and warning events are displayed in yellow ().

The number at the top right of the chart in the gray bar is the average value from the last 72-hour period. Numbers shown at the bottom and top of the trend line graph are the minimum and maximum values for the last 72-hour period. The gray bar below the chart contains the count of active (new and acknowledged) events and obsolete events from the last 72-hour period.



Latency counter chart

The Latency counter chart provides a high-level overview of the object latency for the preceding 72-hour period. Latency refers to the average response time for all I/O requests; expressed in milliseconds per operation, the service time, wait time, or both experienced by a data packet or block in the cluster storage component under consideration.

Top (counter value): The number in the header displays the average for the preceding 72-hour period.

Middle (performance graph): The number at the bottom of the graph displays the lowest latency, and the number at the top of the graph displays the highest latency for the preceding 72-hour period. Position your cursor over the graph trend line to view the latency value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

IOPS counter chart

The IOPS counter chart provides a high-level overview of the object IOPS health for the preceding 72-hour period. IOPS indicates the speed of the storage system in number of input/output operations per second.

Top (counter value): The number in the header displays the average for the preceding 72-hour period.

Middle (performance graph): The number at the bottom of the graph displays the lowest IOPS, and the number at the top of the graph displays the highest IOPS for the preceding 72-hour period. Position your cursor over the graph trend line to view the IOPS value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

MBps counter chart

The MBps counter chart displays the object MBps performance, and indicates how much data has been transferred to and from the object in megabytes per second. The MBps counter chart provides a high-level overview of the object's MBps health for the preceding 72-hour period.

Top (counter value): The number in the header displays the average number of MBps for the preceding 72-hour period.

Middle (performance graph): The value at the bottom of the graph displays the lowest number of MBps, and the value at the top of the graph displays the highest number of MBps for the preceding 72-hour period. Position your cursor over the graph trend line to view the MBps value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

Performance Capacity Used counter chart

The Performance Capacity Used counter chart displays the percentage of performance capacity that is being consumed by the object.

Top (counter value): The number in the header displays the average used performance capacity for the preceding 72-hour period.

Middle (performance graph): The value at the bottom of the graph displays the lowest used performance capacity percentage, and the value at the top of the graph displays the highest used performance capacity percentage for the preceding 72-hour period. Position your cursor over the graph trend line to view the used performance capacity value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

Utilization counter chart

The Utilization counter chart displays the object utilization percentage. The Utilization counter chart provides a high-level overview of the percentage of the object or bandwidth utilization for the preceding 72-hour period.

Top (counter value): The number in the header displays the average utilization percentage for the preceding 72-hour period.

Middle (performance graph): The value at the bottom of the graph displays the lowest utilization percentage, and the value at the top of the graph displays the highest utilization percentage for the preceding 72-hour period. Position your cursor over the graph trend line to view the utilization value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

Events

The events history table, where applicable, lists the most recent events that occurred on that object. Clicking the event name displays details of the event on the Event Details page.

Components of the Performance Explorer page

The Performance Explorer page enables you to compare the performance of similar objects in a cluster—for example, all the volumes in a cluster. This is beneficial when troubleshooting performance events and fine-tuning object performance. You can also compare objects with the root object, which is the baseline against which other object comparisons are made.

You can click the **Favorites** button () to add this object to your list of favorite storage objects. A blue button () indicates that this object is already a favorite.

You can click the **Switch to Health View** button to display the Health details page for this object. In some cases you can learn important information about the storage configuration settings for this object that may help when troubleshooting an issue.

The Performance Explorer page displays a list of cluster objects and their performance data. This page displays all the cluster objects of the same type (for example, volumes and their object-specific performance statistics) in a tabular format. This view provides an efficient overview of cluster object performance.



If "N/A" appears in any cell of the table, it means that a value for that counter is not available because there is no I/O on that object at this time.

The Performance Explorer page contains the following components:

Time Range

Enables you to select a time range for the object data.

You can choose a predefined range, or specify your own custom time range.

View and Compare

Enables you to select which type of correlated object is displayed in the grid.

The options available depend on the root object type and its available data. You can click the View and Compare drop-down list to select an object type. The object type that you select is displayed in the list.

Filtering

Enables you to narrow the amount of data you receive, based on your preferences.

You can create filters that apply to the object data—for example, IOPS greater than 4. You can add up to four simultaneous filters.

Comparing

Displays a list of the objects that you have selected for comparison with the root object.

Data for the objects in the Comparing pane is displayed in the Counter Charts.

View Statistics In

For volume and LUNs, enables you to select whether the statistics are displayed after each collection cycle (default 5 minutes), or whether the statistics are shown as an hourly average. This functionality enables

you to view the latency chart in support of the NetApp "Performance Guarantee" program.

Counter Charts

Displays graphed data for each object performance category.

Typically, only three or four charts are displayed by default. The Choose charts component enables you to display additional charts, or hide specific charts. You can also choose to show or hide the Events Timeline.

Events Timeline

Displays performance and health events occurring across the timeline that you selected in the Time Range component.

Descriptions of the Performance Explorer pages

You use the Performance Explorer pages to view detailed performance information about each of the available storage object; such as clusters, aggregates, volumes, and so on. These pages enable you to assess the overall performance of all objects and compare object performance data in a side-by-side format.

Performance/Cluster Explorer page

The Performance/Cluster Explorer page provides a detailed performance overview of all the clusters that are managed by Unified Manager.

The Performance/Cluster Explorer page enables you to track cluster performance and compare the objects within that cluster during a specific time period, which helps in troubleshooting and fine-tuning the performance of a cluster.

Using the View and Compare functionality you can compare the performance of the cluster with:

- · the nodes on this cluster
- the storage virtual machines (SVMs) of this cluster
- · the aggregates on this cluster

The Performance/Cluster Explorer page enables you to:

- · View threshold-related issues and their details
- · Track cluster performance data
- · Investigate and troubleshoot threshold-related issues
- Investigate and troubleshoot performance issues

Performance/Node Explorer page

The Performance/Node Explorer page provides a detailed performance overview of all nodes within a cluster.

The Performance/Node Explorer page enables you to track and compare node performance during a specific time period, which helps you to troubleshoot and fine-tune the performance of your nodes.

Using the View and Compare functionality you can compare the performance of this node with:

- · other nodes on the same cluster
- the aggregates on the node
- · the ports on the node

The Performance/Node Explorer page enables you to:

- · View threshold-related issues and their details
- · Track and compare node performance data
- Investigate and troubleshoot threshold-related issues
- · Investigate and troubleshoot performance issues

Performance/Aggregate Explorer page

The Performance/Aggregate Explorer page provides a detailed performance overview of all the aggregates in a cluster.

The Performance/Aggregate Explorer page enables you to track and compare aggregate performance during a specific time period, which helps in troubleshooting and fine-tuning the performance of an aggregate.



Root aggregates are not displayed on this page.

Using the View and Compare functionality you can compare the performance of this aggregate with:

- · other aggregates on the same node
- · other aggregates on the same cluster
- · the node on which the aggregate resides
- · all nodes on the cluster that is using this aggregate
- · the volumes that reside on this aggregate

The Performance/Aggregate Explorer page enables you to:

- · View threshold-related issues and their details
- · Track and compare aggregate performance data
- · Investigate and troubleshoot threshold-related issues
- · Investigate and troubleshoot performance issues

Performance/Volume or Performance/FlexGroup Explorer page

This page provides detailed performance information for a volume in a cluster. The title of this page depends on whether you are viewing a FlexVol volume or a FlexGroup volume.

The Volume or FlexGroup Explorer page enables you to track and compare volume performance during a specific time period, which helps you to troubleshoot and fine-tune your volume performance.



Root volumes are not displayed on this page.

Using the View and Compare functionality:

- For FlexVol volumes, you can compare the performance of this volume with:
 - other volumes on the same aggregate
 - the aggregate on which this volume resides
 - the SVM on which this volume resides
 - the LUNs that are on this volume
- For FlexGroup volumes, you can compare the performance of this FlexGroup with:
 - the aggregates on which the FlexGroup resides
 - the SVM on which the FlexGroup resides
 - the constituent volumes of the FlexGroup

The statistics in the charts are updated after each collection period; which by default is every 5 minutes. The View statistics in selector provides an option to show statistics averaged over the previous hour. This functionality enables you to view the latency chart in support of the NetApp "Performance Guarantee" program.

The Performance/Volume Explorer or Performance/FlexGroup Explorer page enables you to:

- · View threshold-related issues and their details
- · Track and compare volume performance data
- · Investigate and troubleshoot threshold-related issues
- · Investigate and troubleshoot performance issues
- Launch System Manager to make a configuration change to the volume

The **Configure Volume** button is available if you are logged in to Unified Manager with the OnCommand Administrator or Storage Administrator role, and when using ONTAP 9.5 or greater.



For data protection (DP) volumes, only counter values for user-generated traffic are displayed. When the cluster is installed with a version of ONTAP software prior to 8.3, no counter values are displayed.

Performance/Constituent Volume Explorer page

The Performance/Constituent Volume Explorer page provides detailed performance information for the selected FlexGroup constituent.

The Performance/Constituent Volume Explorer page enables you to track and compare constituent performance during a specific time period, which helps in troubleshooting and fine-tuning the performance of a FlexGroup volume and its constituent volumes.

Using the View and Compare functionality you can compare the performance of this constituent volume with:

- · the aggregate on which this constituent volume resides
- the SVM on which this constituent volume resides
- the FlexGroup volume to which the constituent volume belongs
- · other volumes that are on the same aggregate

The Performance/Constituent Volume Explorer page enables you to:

- View threshold-related issues and their details
- Track and compare constituent performance data
- · Investigate and troubleshoot threshold-related issues
- Investigate and troubleshoot performance issues



For data protection (DP) volumes, only counter values for user-generated traffic are displayed. When the cluster is installed with a version of ONTAP software prior to 8.3, no counter values are displayed.

Performance/Port Explorer page

The Performance/Port Explorer page provides a detailed performance overview of all ports in a cluster.



Performance counter values are displayed for physical ports only. Counter values are not displayed for VLANs or interface groups.

The Performance/Port Explorer page enables you to track and compare port performance during a specific time period, which helps you to troubleshoot and fine-tune your port performance.

Using the View and Compare functionality you can compare the performance of this port with:

- · other ports on the same node
- · the node on which the port resides
- · LIFs that are on the port



Only cluster and data LIFs are displayed when filtering using the "LIFs on this port" option. No intercluster LIFs are shown.

The Performance/Port Explorer page enables you to:

- View threshold-related issues and their details
- · Track and compare port performance data
- · Investigate and troubleshoot threshold-related issues
- Investigate and troubleshoot performance issues

Performance/SVM Explorer page

The Performance/SVM Explorer page provides a detailed performance overview of all the storage virtual machines (SVMs) in a cluster.

The Performance/SVM Explorer page enables you to track and compare SVM performance during a specific time period, which helps you to troubleshoot and fine-tune your SVM performance.

Using the View and Compare functionality you can compare the performance of this SVM with:

- other SVMs on the same cluster
- · the volumes on this SVM
- · the LIFs on this SVM

The Performance/SVM Explorer page enables you to:

- View threshold-related issues and their details
- · Track and compare SVM performance data
- Investigate and troubleshoot threshold-related issues
- · Investigate and troubleshoot performance issues

Performance/LUN Explorer page

The Performance/LUN Explorer page provides a detailed overview of the performance of all the LUNs within a cluster.

The Performance/LUN Explorer page enables you to track and compare LUN performance during a specific time period, which helps you to troubleshoot and fine-tune the performance of your LUNs.

Using the View and Compare functionality you can compare the performance of this LUN with:

- · other LUNs that are on the same volume
- · the volume on which the LUN resides

The statistics in the charts are updated after each collection period; which by default is every 5 minutes. The View statistics in selector provides an option to show statistics averaged over the previous hour. This functionality enables you to view the latency chart in support of the NetApp "Performance Guarantee" program.

The Performance/LUN Explorer page enables you to:

- View threshold-related issues and their details
- · Track and compare LUN performance data
- Investigate and troubleshoot threshold-related issues
- · Investigate and troubleshoot performance issues

Namespace Explorer page

The Namespace Explorer page provides a detailed overview of the performance of all the Namespaces within a cluster.

The Namespace Explorer page enables you to track and compare Namespace performance during a specific time period, which helps you to troubleshoot and fine-tune the performance of your Namespaces.

Using the View and Compare functionality you can compare the performance of this Namespace with:

- the volume on which the Namespace resides
- other Namespaces that are on the same volume
- other Namespaces that are on the same SVM

The Namespace Explorer page enables you to:

- · View threshold-related issues and their details
- Track and compare Namespace performance data
- Investigate and troubleshoot threshold-related issues
- Investigate and troubleshoot performance issues
- Launch System Manager to make a configuration change to the Namespace

The **Configure Namespace** button is available if you are logged in to Unified Manager with the OnCommand Administrator or Storage Administrator role, and when using ONTAP 9.5 or greater.

Performance/LIF Explorer page

The Performance/LIF Explorer page provides a detailed performance overview for all of the LIFs within a cluster.

The Performance/LIF Explorer page enables you to track and compare LIF performance during a specific time period, which helps you to troubleshoot and fine-tune your LIF performance.

Using the View and Compare functionality you can compare the performance of this LIF with:

- · other LIFs that are on the same port
- · other LIFs that are on the same SVM
- the port on which the LIF resides
- the SVM on which the LIF resides

The Performance/LIF Explorer page enables you to:

- · View threshold-related issues and their details
- · Track and compare LIF performance data
- Investigate and troubleshoot threshold-related issues
- · Investigate and troubleshoot performance issues

Descriptions of the counter charts

You use the Performance Explorer counter charts to view and compare performance data for selected storage objects. These charts can help you to understand performance trends and isolate and resolve performance issues.

Latency performance counter charts

The Performance Explorer Latency counter charts display the number of milliseconds that are required for the selected storage object to respond to application requests.

The popup window that is displayed when your cursor is in the chart area shows the specific counter values at specific times.

The bottom of the chart page displays information for the minimum, maximum, average, and 95th percentile

latency for the selected time range.

There are three types of Latency charts available:

Latency - Total counter chart

Displays the number of milliseconds required to respond to application requests. The average latency values are I/O weighted.

Latency - Breakdown counter chart

Displays the same latency data separated into read, write, and other latency.

This chart option applies when the selected object is an SVM, node, aggregate, volume, or LUN.

Latency - Cluster Components counter chart

Displays the latency data by cluster component. This enables you to identify the cluster component that is responsible for the latency. By hovering your cursor in the chart you can view the exact latency contribution for each component.

This chart option applies when the selected object is an SVM, node, aggregate, volume, or LUN.

Zoom View button

Displays a magnified view of the counter chart data.

Events

The occurrence of critical, warning, and informational events are indicated on the time lines above the charts.

Thresholds

The dashed, horizontal line indicates the utilization warning threshold value set in Unified Manager.

The solid red line indicates the utilization critical threshold value set in Unified Manager.

Counters

The counters in the left pane show which counter values are being displayed. Deselecting or selecting the that is associated with a counter hides and shows that counter information from the chart and can help when comparing object latency.

IOPS performance counter charts

The Performance Explorer IOPS counter charts display the number of input/output operations processed per second by the selected storage object.

The popup window that is displayed when your cursor is in the chart area shows the specific counter values at specific times.

When displayed in Zoom view, the volume and LUN IOPS charts also display Quality of Service (QoS) maximum and minimum throughput threshold settings, if configured in ONTAP. The IOPS/TB charts display QoS peak and expected throughput threshold settings, if configured.

When viewing a volume or LUN that is sharing the IOPS of a shared QoS policy, a line for "Total Workload IOPS" is displayed to show the IOPS that are being used by all other workloads sharing this policy.

The bottom of the chart page displays information for the minimum, maximum, average, and 95th percentile IOPS for the selected time range.

There are four types of IOPS charts available:

IOPS - Total counter chart

Displays the number of input/output operations processed per second.

IOPS - Breakdown counter chart

Displays the same IOPS data separated into read, write, and other IOPS.

This chart option applies when the selected object is an SVM, node, aggregate, volume, or LUN.

IOPS - Protocols counter chart

Displays the same IOPS data, but for SVMs the performance data is separated into individual components for CIFS, NFS, FCP, NVMe, and iSCSI protocol traffic.

IOPS/TB - Total counter chart

Displays the number of input/output operations processed per second based on the total logical space that is being consumed by the volume, in terabytes. Also called I/O density, this counter measures how much performance can be delivered by a given amount of storage capacity.

This chart option is available only when the selected object is a volume. It displays performance data only when the logical capacity used by the volume is greater than or equal to 1 TB. Gaps will be displayed in the chart when the used capacity falls below 1 TB during the selected timeframe.



In some situations the IOPS/TB chart might show large spikes in data while the IOPS chart does not show the same behavior. This is a known limitation where some performance data is captured less frequently than other performance data. This chart will typically return to normal operation in 5 or 10 minutes when the collection cycles are synchronized.

Zoom View button

Displays a magnified view of the counter chart data.

Events

The occurrence of critical, error, warning, and informational events are indicated on the time lines above the charts.

• Thresholds

The dashed, horizontal line indicates the utilization warning threshold value set in Unified Manager.

The solid red line indicates the utilization critical threshold value set in Unified Manager.

Counters

The counters in the left pane show which counter values are being displayed. Deselecting or selecting the that is associated with a counter hides and shows that counter information from the chart and can help when comparing object IOPS.

MBps performance counter charts

The Performance Explorer MBps counter charts display the number of megabytes of data transferred to and from the selected object per second.

The popup window that is displayed when your cursor is in the chart area shows the specific counter values at specific times.

When displayed in Zoom view, the volume and LUN charts also display Quality of Service (QoS) maximum MBps throughput threshold settings, if configured in ONTAP.

When viewing a volume or LUN that is sharing the MBps of a shared QoS policy, a line for "Total Workload MBps" is displayed to show the MBps that are being used by all other workloads sharing this policy.

The bottom of the chart page displays information for the minimum, maximum, average, and 95th percentile MBps for the selected time range.

There are two types of MBps charts available:

MBps - Total counter chart

Displays the number of megabytes of data transferred to and from the selected object per second.

MBps - Breakdown counter chart

Displays the same MBps data separated into disk read, Flash Cache read, write, and other operations.

This chart option applies when the selected object is an SVM, node, aggregate, volume, or LUN.



Flash Cache data is displayed only for nodes, and only when a Flash Cache module is installed in the node.

Zoom View button

Displays a magnified view of the counter chart data.

Events

The occurrence of critical, error, warning, and informational events are indicated on the time lines above the charts.

Thresholds

The dashed, horizontal line indicates the utilization warning threshold value set in Unified Manager.

The solid red line indicates the utilization critical threshold value set in Unified Manager.

Counters

The counters in the left pane show which counter values are being displayed. Deselecting or selecting the

• that is associated with a counter hides and shows that counter information from the chart and can help when comparing object MBps.

Utilization performance counter chart

The Performance Explorer Utilization counter chart displays the average percentage of the selected resource that is being used.

The popup window that is displayed when your cursor is in the chart area shows the specific counter values at specific times.

The bottom of the chart page displays information for the minimum, maximum, average, and 95th percentile utilization for the selected time range.

Utilization - Total counter chart

Displays the average percentage of the selected resource that is being used. For nodes this indicates utilization of node resources (CPU and RAM), for aggregates this indicates utilization of the disks in the aggregate, and for ports this indicates the bandwidth utilization of the port.

This chart option applies when the selected object is a node, aggregate, or port.

Zoom View button

Displays a magnified view of the counter chart data.

Events

The occurrence of critical, warning, and informational events are indicated on the time lines above the charts.

Thresholds

The dashed, horizontal line indicates the utilization warning threshold value set in Unified Manager.

The solid red line indicates the utilization critical threshold value set in Unified Manager.

Counters

The counters in the left pane show which counter values are being displayed. Deselecting or selecting the that is associated with a counter hides and shows that counter information from the chart and can help when comparing object utilization.

Performance Capacity Used performance counter charts

The Performance ExplorerPerformance Capacity Used counter charts display the percentage of performance capacity that is being consumed by the node or aggregate.



Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

These charts apply only when the selected object is a node or aggregate.

The popup window that is displayed when your cursor is in the chart area shows the specific counter values at specific times.

The bottom of the chart page displays information for the minimum, maximum, average, and 95th percentile performance capacity used for the selected time range.

There are two types of Performance Capacity Used charts available:

Performance Capacity Used - Total counter chart

Displays the percentage of performance capacity that is being consumed by the node or aggregate.

Green zone

The capacity value is under the warning threshold set in Unified Manager.

Yellow zone

The capacity value is approaching the warning threshold set in Unified Manager.

· Red zone

The capacity value is above the warning threshold and approaching the maximum threshold set in Unified Manager.

Performance Capacity Used - Breakdown counter chart

Displays the same percentage of performance capacity separated into user protocols, system background processes, and the amount of free performance capacity.

Zoom View button

Displays a magnified view of the counter chart data.

Events

The occurrence of critical, warning, and informational events are indicated on the time lines above the charts.

Thresholds

The dashed, horizontal line indicates the capacity warning threshold value set in Unified Manager.

The solid red line indicates the capacity critical threshold value set in Unified Manager.

The solid black line at 100% is the recommended maximum performance capacity used value.

Counters

The counters in the left pane show which counter values are being displayed. Deselecting or selecting the that is associated with a counter can restrict the comparison.

Available IOPS performance counter chart

The Performance Explorer Available IOPS counter chart displays the number of

input/output operations per second that are currently available (free) on the selected storage object.

The popup window that is displayed when your cursor is in the chart area shows the specific counter values at specific times.

This chart option applies only when the selected object is a node or aggregate.

The bottom of the chart page displays information for the minimum, maximum, average, and 95th percentile performance capacity used for the selected time range.

Available IOPS - Total counter chart

Displays the number of input/output operations per second that are currently available (free) on the selected storage object. This number is the result of subtracting the currently used IOPS from the total IOPS that Unified Manager calculates that the object can perform.



Available IOPS data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

Zoom View button

Displays a magnified view of the counter chart data.

Events

The occurrence of critical, warning, and informational events are indicated on the time lines above the charts.

Counters

The counters in the left pane show which counter values are being displayed. Deselecting or selecting the that is associated with a counter hides and shows that counter information from the chart and can help when comparing objects.

Cache Miss Ratio performance counter chart

The Performance Explorer Cache Miss Ratio counter chart displays the percentage of read requests from client applications that are returned from the disk instead of being returned from the cache.

The popup window that is displayed when your cursor is in the chart area shows the specific counter values at specific times.

The bottom of the chart page displays information for the minimum, maximum, average, and 95th percentile cache miss ratio for the selected time range.

Cache Miss Ratio - Total counter chart

Displays the percentage of read requests from client applications that are returned from the disk instead of being returned from the cache.

This chart option applies only when the selected object is a volume.

Zoom View button

Displays a magnified view of the counter chart data.

Events

The occurrence of critical, warning, and informational events are indicated on the time lines above the charts.

Counters

The counters in the left pane show which counter values are being displayed. Deselecting or selecting the that is associated with a counter hides and shows that counter information from the chart and can help when comparing objects.

Viewing object configuration information

The object Information pages—located on the landing page of each object—display the values for the non-performance configuration attributes of each storage object. Some of the attributes are physical configuration settings, while other attributes may affect the performance of the object.

For example, it is useful to know the amount of space that is available for an aggregate or for a node. Knowing the speed setting for a particular port can help when you are diagnosing a performance issue.

Performance/Cluster Information page

Use the Performance/Cluster Information page to view a list of the physical and logical attributes of the cluster. This information might help in answering performance-related questions.

Cluster attributes

Management LIF

The name of the cluster management LIF, and whether the LIF is currently available (Up), or not (Down).

IP Address

The IPv4 or IPv6 address of the cluster management LIF.

FQDN

The fully qualified domain name (FQDN) of the cluster management LIF.

OS Version

The version of ONTAP software installed on the cluster.



If different versions of ONTAP software are installed on the nodes in the cluster, the listed version is the lowest version number. Check the Performance/Node Information page to view the version of ONTAP software installed on each node.

Serial Number

The unique identification number of the cluster.

· Model / Family

The platform model number and model family of all the nodes in the cluster.

Capacity (free/total)

The total storage available to the cluster, in gigabytes, and the amount of storage currently available.

Allowed Protocols

The list of all protocols that can be serviced by this cluster. The available protocols are FC/FCoE, iSCSI, HTTP, NVMe, NDMP, NFS, and CIFS.

Nodes

The number of nodes in this cluster. You can click the number to display the nodes in the Performance/Node Inventory page.

Storage Virtual Machines

The number of SVMs in this cluster. You can click the number to display the SVMs in the Performance/SVM Inventory page.

• LIFs

The number of LIFs in this cluster. You can click the number to display the LIFs in the Performance/LIF Inventory page.

Contact / Location

If available, the name of the storage administrator to contact regarding this cluster, and the location of the cluster.

Performance/Node Information page

Use the Performance/Node Information page to view a list of the physical and logical attributes of the node. This information might help in answering performance-related questions.

Node attributes

IP Address

The IPv4 or IPv6 address of the node management LIF.

FQDN

The fully qualified domain name (FQDN) of the node management LIF.

OS Version

The version of ONTAP software installed on the node.

· Model / Family

The platform model number of the node.

Capacity (free/total)

The total storage available to the node, in gigabytes, and the amount of storage currently available.

Cluster

The name of the cluster to which this node belongs. You can click the name to display cluster details the Performance/Cluster Explorer page.

HA Partner

The name of the HA partner node, if applicable. You can click the name to display partner node details in the Performance/Node Explorer page.

Aggregates

The number of aggregates on this node. You can click the number to display the aggregates in the Performance/Aggregates Inventory page.



The number listed here may not match the number in the Performance/Aggregates Inventory page because the inventory page does not include root aggregates.

Ports

The number of ports on this node. You can click the number to display the ports in the Performance/Ports Inventory page.



The number listed here may not match the number in the Performance/Ports Inventory page because the inventory page does not include node management ports.

Contact / Location

If available, the name of the administrator to contact regarding this node, and the location of the node.

• # of Cores / Speed

If available, the number of CPU cores on the controller, and the speed of the CPU cores.

RAM

If available, the total memory available on the controller.

Flash Devices



Flash Cache data is displayed only for nodes, and only when a Flash Cache module is installed in the node.

Slot Number

The slot number in which the Flash Cache module is installed.

Status

The operational status of the module. Valid values:

- · Online
- · Offline failed
- · Offline_threshold

· Model / Family

The model number of the module.

Firmware Rev

The version of firmware installed on the module.

Capacity

The size of the installed Flash Cache module.

Performance/Aggregate Information page

Use the Performance/Aggregate Information page to view a list of the physical and logical attributes of the aggregate. This information might help in answering performance-related questions.

Aggregate attributes

· Aggregate Type

The type of aggregate:

- HDD
- Hybrid

Combines HDDs and SSDs, but Flash Pool has not been enabled.

Hybrid (Flash Pool)

Combines HDDs and SSDs, and Flash Pool has been enabled.

- SSD
- SSD (FabricPool)

Combines SSDs and a cloud tier

VMDisk (SDS)

Virtual disks within a virtual machine

VMDisk (FabricPool)

Combines virtual disks and a cloud tier

LUN (FlexArray)

Cluster

The name of the cluster to which the aggregate belongs. You can click the name to display cluster details in the Performance/Cluster Explorer page.

Node

The name of the node to which the disks of the aggregate belong. You can click the name to display node details in the Performance/Node Explorer page.

Flash Pool

Whether this is a Flash Pool aggregate: Yes or No.

A Flash Pool aggregate is a hybrid aggregate that consists of both SSDs and HDDs.

FabricPool

Whether this is a FabricPool aggregate: Yes or No.

A FabricPool aggregate is an aggregate that consists of both SSDs and a cloud tier.

Inactive Data Reporting

Whether the inactive data reporting capability is enabled or disabled on this aggregate. When enabled, volumes on this aggregate display the amount of cold data in the Performance/Volumes inventory page.

The value in this field is "N/A" when the version of ONTAP does not support inactive data reporting.

Performance/Volume or Performance/FlexGroup Information page

Use this page to view a list of the physical and logical attributes of the volume. This information might help in answering performance-related questions. The title of this page depends on whether you are viewing a FlexVol volume or a FlexGroup volume.

Volume attributes

Type

The volume's type; either read-write (RW) or data-protection (DP).

Style

The style of volume; either FlexVol or FlexGroup.



The performance pages of Unified Manager do not support Infinite Volumes.

Cluster

The name of the cluster to which this FlexVol volume or FlexGroup volume belongs. You can click the name to display cluster details in the Performance/Cluster Explorer page.

Aggregates

The name of the aggregate on which this FlexVol volume resides, or the number of aggregates on which this FlexGroup volume resides.

For FlexVol volumes, you can click the name to display aggregate details in the Performance/Aggregate Explorer page. For FlexGroup volumes, you can click the number to display the aggregates that are used in this FlexGroup volume in the Performance/Aggregate Inventory page.

Storage Virtual Machine

The name of the SVM to which this FlexVol volume or FlexGroup volume belongs. You can click the name to display SVM details in the Performance/SVM Explorer page.

Tiering Policy

The tiering policy set on the volume. The policy takes affect only when the volume is deployed on a FabricPool aggregate. The available policies are:

- None. The data for this volume always remains on the performance tier.
- Snapshot Only. Only Snapshot data is moved automatically to the cloud tier. All other data remains on the performance tier.
- Backup. On data protection volumes, all transferred user data starts in the cloud tier, but later client reads can cause hot data to move to the performance tier.
- Auto. Data on this volume is moved between the performance tier and the cloud tier automatically when ONTAP determines that the data is "hot" or "cold".

RAID Type

The redundancy type that is being used on the performance tier of the aggregate where this volume resides. Possible types:

- RAID0
- RAID4
- RAID-DP
- RAID-TEC



The value "Not Applicable" is displayed for FlexGroup volumes because the constituent volumes can be on aggregates of different RAID types.

Capacity (free/total)

The total storage available on the volume, in gigabytes, and the amount of storage currently available.

Performance/Constituent Volume Information page

Use the Performance/Constituent Volume Information page to view a list of the physical and logical attributes of the FlexGroup constituent volume. This information might help in

answering performance-related questions.

Constituent Volume attributes

Type

The constituent's type; either read-write (RW) or data-protection (DP).

Style

The style of volume; this is a constituent volume of a FlexGroup volume.

Cluster

The name of the cluster to which this FlexGroup constituent volume belongs. You can click the name to display cluster details in the Performance/Cluster Explorer page.

Aggregate

The name of the aggregate on which this FlexGroup constituent volume resides. You can click the name to display aggregate details in the Performance/Aggregate Explorer page.

FlexGroup

The name of the FlexGroup volume to which this constituent belongs. You can click the name to display FlexGroup volume details in the Performance/FlexGroup Explorer page.

Storage Virtual Machine

The name of the SVM to which this FlexGroup constituent volume belongs. You can click the name to display SVM details in the Performance/SVM Explorer page.

Tiering Policy

The tiering policy set on the volume. The policy takes affect only when the volume is deployed on a FabricPool aggregate. The available policies are:

- None. The data for this volume always remains on the performance tier.
- Snapshot Only. Only Snapshot data is moved automatically to the cloud tier. All other data remains on the performance tier.
- Backup. On data protection volumes, all transferred user data starts in the cloud tier, but later client reads can cause hot data to move to the performance tier.
- Auto. Data on this volume is moved between the performance tier and the cloud tier automatically when ONTAP determines that the data is "hot" or "cold".

RAID Type

The redundancy type that is being used on the aggregate where this constituent resides. Possible types:

- RAID0
- RAID4
- RAID-DP
- · RAID-TEC

Capacity (free/total)

The total storage available on the constituent, in gigabytes, and the amount of storage currently available.

Performance/Port Information page

Use the Performance/Port Information page to view a list of the physical and logical attributes of the port. This information might help in answering performance-related questions.

Port attributes

• WWN

The WWN (World Wide Name) of the port.

Node

The name of the node on which the physical port resides. You can click the name to display node details in the Performance/Node Explorer page.

Cluster

The name of the cluster to which the port belongs. You can click the name to display cluster details the Performance/Cluster Explorer page.

· Operational Speed

The actual speed at which the port is configured to run.

FCP ports are auto-sensing and display as "Auto".

Role

The network port function: either Data or Cluster.

FCP ports cannot have a role, and this field is not displayed.

Type

The port type: either Network or FCP (Fibre Channel Protocol).

State

The link status of the port.

- For network ports, an active port is listed as "Up" and an inactive port is listed as "Down".
- For FCP ports, an active port is listed as "Online" and an inactive port is listed as "Link not connected".

Performance/SVM Information page

Use the Performance/SVM Information page to view a list of the configured attributes of the SVM. This information might help in answering performance-related questions.

SVM attributes

IP Address

If defined, this is the IPv4 or IPv6 address of the SVM management LIF.

IPspace

The IPspace in which this SVM resides.

Domain Name

The fully qualified domain name (FQDN) of the SVM management LIF.

Service Type

The type of SVM.

Possible values include: "Admin" for the cluster-wide management SVM, "System" for cluster-level communications in an IPspace, "Data" for data serving SVM, and "Node" for node management SVM.

Capacity (free/total)

The total storage available to the SVM, in gigabytes, and the amount of storage currently available.

Cluster

The name of the cluster to which the SVM belongs. You can click the name to display cluster details in the Performance/Cluster Explorer page.

Volumes

The number of volumes in the SVM. You can click the number to display the volumes in the Performance/Volume Inventory page.

LIFs

The number of LIFs available to the SVM. You can click the number to display the LIFs in the Performance/LIFs Inventory page.

Data LIFs

The number and type of Data LIFs available to the SVM.

Allowed Volume Type

The type of volume that can be created on the SVM.

SVMs can contain one or more FlexVol volumes or FlexGroup volumes. The FlexGroup type is allowed when using ONTAP 9.1 or later.



The performance pages of Unified Manager do not support Infinite Volumes.

Allowed Protocols

The list of all protocols that can be serviced by this SVM. The available protocols are FC/FCoE, iSCSI,

HTTP, NDMP, NVMe, NFS, and CIFS.

Port Set

If defined for FCP or iSCSI protocols, the port set that is assigned to this SVM.

Performance/LUN Information page

Use the Performance/LUN Information page to view a list of the physical and logical attributes of the LUN. This information might help in answering performance-related questions.

LUN attributes

• WWN

The WWN (World Wide Name) of the LUN.

Path

The full path of the LUN, for example, /vol/vol1/lun1.

Alignment

Indicates the alignment state of the LUN. Possible values:

- Not mapped
- · Aligned
- Misaligned
- · Possibly misaligned
- Indeterminate

Capacity (free/total)

The total storage available on the LUN, in gigabytes, and the amount of storage currently available.

Volume

The name of the volume to which the LUN belongs. You can click the name to display volume details in the Performance/Volume Explorer page.

Storage Virtual Machine

The name of the SVM to which the LUN belongs. You can click the name to display SVM details in the Performance/SVM Explorer page.

Node

The name of the node on which the LUN resides. You can click the name to display node details in the Performance/Node Explorer page.

Cluster

The name of the cluster to which the LUN belongs. You can click the name to display cluster details in the Performance/Cluster Explorer page.

State

The state of the LUN. Valid states can be online, offline, nvfail, space-error, and foreign-lun-error.

Mapped

Whether the LUN is mapped to an initiator group (true), or not (false).

Namespace Information page

Use the Namespace Information page to view a list of the physical and logical attributes of the Namespace. This information might help in answering performance-related questions.

Namespace attributes

Cluster

The name of the cluster to which the Namespace belongs. You can click the name to display cluster details in the Performance/Cluster Explorer page.

· Capacity (free/total)

The total storage capacity of the Namespace and the amount of storage currently available.

Node

The name of the node on which the Namespace resides. You can click the name to display node details in the Performance/Node Explorer page.

Path

The full path of the Namespace, for example, /vol/vol1/namespace1.

State

The state of the Namespace. Valid states can be online, offline, nvfail, and space-error.

Subsystem

The subsystem of the Namespace.

Storage Virtual Machine

The name of the SVM to which the Namespace belongs. You can click the name to display SVM details in the Performance/SVM Explorer page.

Volume

The name of the volume to which the Namespace belongs. You can click the name to display volume details in the Performance/Volume Explorer page.

Performance/LIF Information page

Use the Performance/LIF Information page to view a list of the configured attributes of the LIF. This information might help in answering performance-related questions.

LIF attributes

IP Address

The IPv4 or IPv6 address assigned to the LIF. There can be multiple IP addresses assigned to a LIF.

Role

The role determines the kind of traffic that is supported over the LIF.

LIFs can have one of the following roles:

- Data
- Cluster
- Node Management
- Intercluster

Failover Group

The name of the failover group that is assigned to the LIF.

This field applies only to network LIFs, not to SAN (FC/ISCSI) and NVMe LIFs.

Failover Policy

The name of the failover policy that is assigned to the LIF.

This field applies only to network LIFs, not to SAN (FC/ISCSI) and NVMe LIFs.

Home Port

The name of the node and port that has been defined as the home port for this interface. You can click the name to display port details in the Performance/Port Explorer page.

Current Port

The name of the node and port on which the interface is currently hosted. You can click the name to display port details in the Performance/Port Explorer page.

Understanding and using the Node Failover Planning page

The Performance/Node Failover Planning page estimates the performance impact on a node if the node's high-availability (HA) partner node fails. Unified Manager bases the estimates on the historical performance of the nodes in the HA pair.

Estimating the performance impact of a failover helps you to plan in the following scenarios:

- If a failover consistently degrades the takeover node's estimated performance to an unacceptable level, you can consider taking corrective actions to reduce the performance impact due to a failover.
- Before initiating a manual failover to perform hardware maintenance tasks, you can estimate how the failover affects the performance of the takeover node in order to determine the best time to perform the task.

Using the Node Failover Planning page to determine corrective actions

Based on the information that is displayed in the Performance/Node Failover Planning page, you can take actions to ensure that a failover does not cause the performance of an HA pair to drop below an acceptable level.

For example, to reduce the estimated performance impact of a failover, you can move some volumes or LUNs from a node in the HA pair to other nodes in the cluster. Doing so ensures that the primary node can continue to deliver acceptable performance after a failover.

Components of the Node Failover Planning page

The components of the Performance/Node Failover Planning page are displayed in a grid and in the Comparing pane. These sections enable you to assess the impact of a node failover on the performance of the takeover node.

Performance statistics grid

The Performance/Node Failover Planning page displays a grid containing statistics for latency, IOPS, utilization, and performance capacity used.



IOPS values displayed in this page and in the Performance/Node Performance Explorer page might not be the same.

In the grid, each node is assigned one of the following roles:

Primary

The node that takes over for the HA partner when the partner fails. The root object is always the Primary node.

Partner

The node that fails in the failover scenario.

· Estimated Takeover

The same as the Primary node. Performance statistics displayed for this node show the takeover node's performance after it takes over the failed partner.



Although the workload of the takeover node is equivalent to the combined workloads of both nodes after a failover, the statistics for the Estimated Takeover node are not the sum of the statistics of the Primary node and the Partner node. For example, if the latency of the Primary node is 2 ms/op and the latency of the Partner node is 3 ms/op, the Estimated Takeover node might have a latency of 4 ms/op. This value is a calculation that Unified Manager performs.

You can click the name of the Partner node if you want it to become the root object. After the Performance/Node Performance Explorer page is displayed, you can click the **Failover Planning** tab to see how performance changes in this node failure scenario. For example, if Node1 is the Primary node and Node2 is the Partner node, you can click Node2 to make it the Primary node. In this way, you can see how the estimated performance changes depending on which node fails.

Comparing pane

The following list describes the components displayed in the Comparing pane by default:

· Events charts

They are displayed in the same format as those in the Performance/Node Performance Explorer page. They pertain to the Primary node only.

Counter charts

They display historical statistics for the performance counter shown in the grid. In each chart, the graph for the Estimated Takeover node shows the estimated performance if a failover had occurred at any given time.

For example, suppose the Utilization chart shows 73% for the Estimated Takeover node at 11 a.m. on February 8. If a failover had occurred at that time, the utilization of the takeover node would have been 73%.

The historical statistics help you find the optimal time for initiating a failover, minimizing the possibility of overloading the takeover node. You can schedule a failover only at times when the predicted performance of the takeover node is acceptable.

By default, statistics for both the root object and the partner node are displayed in the Comparing pane. Unlike in the Performance/Node Performance Explorer page, this page does not display the **Add** button for you to add objects for statistics comparison.

You can customize the Comparing pane in the same way as you do in the Performance/Node Performance Explorer page. The following list shows examples of customizing the charts:

- Click a node name to show or hide the node's statistics in the Counter charts.
- Click **Zoom View** to display a detailed chart for a particular counter in a new window.

Using a threshold policy with the Node Failover Planning page

You can create a node threshold policy so that you can be notified in the Performance/Node Failover Planning page when a potential failover would degrade the performance of the takeover node to an unacceptable level.

The system-defined performance threshold policy named "Node HA pair over-utilized" generates a warning event if the threshold is breached for six consecutive collection periods (30 minutes). The threshold is considered breached if the combined performance capacity used of the nodes in an HA pair exceeds 200%.

The event from the system-defined threshold policy alerts you to the fact that a failover will cause the latency of the takeover node to increase to an unacceptable level. When you see an event that is generated by this policy for a particular node, you can navigate to the Performance/Node Failover Planning page for that node to view the predicted latency value due to a failover.

In addition to using this system-defined threshold policy, you can create threshold policies by using the "Performance Capacity Used - Takeover" counter, and then apply the policy to selected nodes. Specifying a threshold lower than 200% enables you to receive an event before the threshold for the system-defined policy is breached. You can also specify the minimum period of time for which the threshold is exceeded to less than 30 minutes if you want to be notified before the system-defined policy event is generated.

For example, you can define a threshold policy to generate a warning event if the combined performance capacity used of the nodes in an HA pair exceeds 175% for more than 10 minutes. You can apply this policy to Node1 and Node2, which form an HA pair. After receiving a warning event notification for either Node1 or Node2, you can view the Performance/Node Failover Planning page for that node to assess the estimated performance impact on the takeover node. You can take corrective actions to avoid overloading the takeover node if a failover does happen. If you take action when the combined performance capacity used of the nodes is under 200%, the takeover node's latency does not reach an unacceptable level even if a failover happens during this time.

Using the Performance Capacity Used Breakdown chart for failover planning

The detailed Performance Capacity Used - Breakdown chart shows the performance capacity used for the Primary node and the Partner node. It also shows the amount of free performance capacity on the Estimated Takeover node. This information helps you determine whether you might have a performance issue if the partner node fails.

About this task

In addition to showing the total performance capacity used for the nodes, the Breakdown chart breaks the values for each node into user protocols and background processes.

- User protocols are the I/O operations from user applications to and from the cluster.
- Background processes are the internal system processes involved with storage efficiency, data replication, and system health.

This additional level of detail enables you to determine whether a performance issue is caused by user application activity or background system processes, such as deduplication, RAID reconstruct, disk scrubbing, and SnapMirror copies.

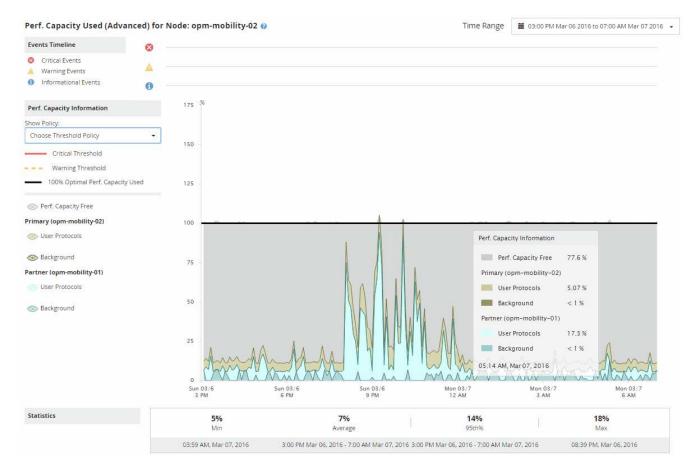
Steps

- 1. Go to the **Performance/Node Failover Planning** page for the node that will serve as the Estimated Takeover node.
- 2. From the **Time Range** selector, choose the period of time for which the historical statistics are displayed in the counter grid and counter charts.

The counter charts with statistics for the Primary node, Partner node, and Estimated Takeover node are displayed.

- 3. From the Choose charts list, select Perf. Capacity Used.
- 4. In the Perf. Capacity Used chart, select Breakdown and click Zoom View.

The detailed chart for Perf. Capacity Used is displayed.



5. Move the cursor over the detailed chart to view the performance capacity used information in the popup window.

The Perf. Capacity Free percentage is the performance capacity available on the Estimated Takeover node. It indicates how much performance capacity is left on the takeover node after a failover. If it is 0%, a failover will cause the latency to increase to an unacceptable level on the takeover node.

6. Consider taking corrective actions to avoid a low performance capacity free percentage.

If you plan to initiate a failover for node maintenance, choose a time to fail the partner node when the performance capacity free percentage is not at 0.

Collecting data and monitoring workload performance

Unified Manager collects and analyzes workload activity every 5 minutes to identify performance events, and it detects configuration changes every 15 minutes. It retains a maximum of 30 days of 5-minute historical performance and event data, and it uses this data to forecast the expected range for all monitored workloads.



This chapter describes how dynamic thresholds work and how they are used to help monitor workload performance. This chapter is not applicable for statistics or events caused by user-defined or system-defined performance threshold breaches.

Unified Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the expected range for I/O response time and operations can be displayed on the Performance/Volume

Details page and in the Event details page. While this activity is being collected, the expected range does not display all changes occurring from workload activity. After collecting 3 days of activity, Unified Manager adjusts the expected range, every 24 hours at 12:00 a.m., to reflect workload activity changes and establish a more accurate performance threshold.

During the first 4 days that Unified Manager is monitoring a volume, if more than 24 hours have passed since the last data collection, the charts on the Performance/Volume Details page will not display the expected range for that volume. Events detected prior to the last collection are still available.



Daylight savings time (DST) changes the system time, which alters the expected range of performance statistics for monitored workloads. Unified Manager immediately begins to correct the expected range, which takes approximately 15 days to complete. During this time you can continue to use Unified Manager, but, since Unified Manager uses the expected range to detect events, some events might not be accurate. Events detected prior to the time change are not affected. Manually changing the time on a cluster, or on a Unified Manager server, to an earlier time will also affect the event analysis results.

Types of workloads monitored by Unified Manager

You can use Unified Manager to monitor the performance of two types of workloads: user-defined and system-defined.

User-defined workloads

The I/O throughput from applications to the cluster. These are processes involved in read and write requests. A FlexVol volume or FlexGroup volume is a user-defined workload.



Unified Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

If one or more of the following is true for a workload, it cannot be monitored by Unified Manager:

- It is a data protection (DP) copy in read-only mode. (Note that when using ONTAP 8.3 and later, DP volumes are monitored for user-generated traffic.)
- It is an Infinite Volume.
- It is an offline data clone.
- It is a mirrored volume in a MetroCluster configuration.

System-defined workloads

The internal processes involved with storage efficiency, data replication, and system health, including:

- Storage efficiency, such as deduplication
- Disk health, which includes RAID reconstruct, disk scrubbing, and so on
- · Data replication, such as SnapMirror copies
- Management activities
- · File system health, which includes various WAFL activities
- File system scanners, such as WAFL scan
- Copy offload, such as offloaded storage efficiency operations from VMware hosts

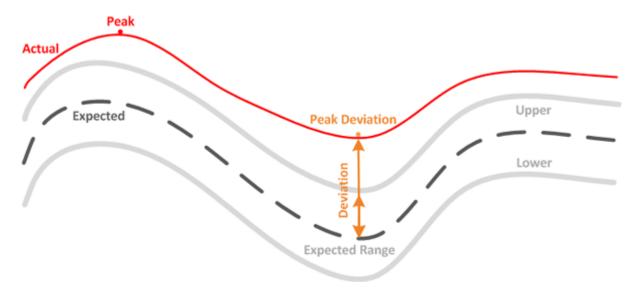
- System health, such as volume moves, data compression, and so on
- Unmonitored volumes

Performance data for system-defined workloads is displayed in the GUI only when the cluster component used by these workloads is in contention. For example, you cannot search for the name of a system-defined workload to view its performance data in the GUI. If multiple system-defined workloads of the same type are displayed, a letter is appended to the workload name. The letter is intended for use by support personnel.

Workload performance measurement values

Unified Manager measures the performance of workloads on a cluster based on historical and expected statistical values, which form the expected range of values for the workloads. It compares the actual workload statistical values to the expected range to determine when workload performance is too high or too low. A workload that is not performing as expected triggers a performance event report to notify you.

In the following illustration, the actual value, in red, represents the actual performance statistics in the time frame. The actual value has crossed the performance threshold, which is the upper bounds of the expected range. The peak is the highest actual value in the time frame. The deviation measures the change between the expected values and the actual values, while the peak deviation indicates the largest change between the expected values and the actual values.



The following table lists the workload performance measurement values.

Measurement	Description
Activity	The percentage of the QoS limit used by the workloads in the policy group.
	If Unified Manager detects a change to a policy group, such as adding or removing a volume or changing the QoS limit, the actual and expected values might exceed 100% of the set limit. If a value exceeds 100% of the set limit it is displayed as >100%. If a value is less than 1% of the set limit it is displayed as <1%.
Actual	The measured performance value at a specific time for a given workload.
Deviation	The change between the expected values and the actual values. It is the ratio of the actual value minus the expected value to the upper value of the expected range minus the expected value.
	A negative deviation value indicates that workload performance is lower than expected, while a positive deviation value indicates that workload performance is higher than expected. If the expected values and the actual value are very low, in the hundredths or thousandths of a percent for example, the deviation will display N/A.
Expected	The expected values are based on the analysis of historical performance data for a given workload. Unified Manager analyzes these statistical values to determine the expected range of values.
Expected Range	The expected range of values is a forecast, or prediction, of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Unified Manager triggers a performance event alert.
Peak	The maximum value measured over a period of time.
Peak Deviation	The maximum deviation value measured over a period of time.

Measurement	Description
Queue Depth	The number of pending I/O requests that are waiting at the interconnect component.
Utilization	For the network processing, data processing, and aggregate components, the percentage of busy time to complete workload operations over a period of time. For example, the percentage of time for the network processing or data processing components to process an I/O request or for an aggregate to fulfill a read or write request.
Write Throughput	The amount of write throughput, in Megabytes per second (MBps), from workloads on a local cluster to the partner cluster in a MetroCluster configuration.

What the expected range of performance is

The expected range of values is a forecast, or prediction, of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Unified Manager triggers a performance event alert.

For example, during regular business hours between 9:00 a.m. and 5:00 p.m., most employees might check their email between 9:00 a.m. and 10:30 a.m. The increased demand on the email servers means an increase in workload activity on the back-end storage during this time. Employees might notice slow response time from their email clients.

During the lunch hour between 12:00 p.m. and 1:00 p.m. and at the end of the work day after 5:00 p.m., most employees are likely away from their computers. The demand on the email servers typically decreases, also decreasing the demand on back-end storage. Alternatively, there could be scheduled workload operations, such as storage backups or virus scanning, that start after 5:00 p.m. and increase activity on the back-end storage.

Over several days, the increase and decrease in workload activity determines the expected range of activity, with upper and lower boundaries for a workload. When the actual workload activity for an object is outside the upper or lower boundaries, and remains outside the boundaries for a period of time, this might indicate that the object is being overused or underused.

How the expected range is formed

Unified Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the expected range for I/O response time and operations can be displayed in the GUI. The minimum required data collection does not account for all changes occurring from workload activity. After collecting the first 3 days of activity, Unified Manager adjusts the expected range, every 24 hours at 12:00 a.m., to reflect workload activity changes and establish a more accurate performance threshold.



Daylight savings time (DST) changes the system time, which alters the expected range of performance statistics for monitored workloads. Unified Manager immediately begins to correct the expected range, which takes approximately 15 days to complete. During this time you can continue to use Unified Manager, but, since Unified Manager uses the expected range to detect events, some events might not be accurate. Events detected prior to the time change are not affected. Manually changing the time on a cluster, or on a Unified Manager server, to an earlier time will also affect the event analysis results.

How the expected range is used in performance analysis

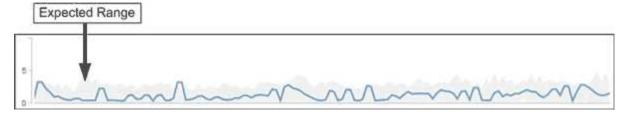
Unified Manager uses the expected range to represent the typical I/O latency (response time) and IOPS (operations) activity for your monitored workloads. It alerts you when the actual latency for a workload is above the upper bounds of the expected range, which triggers a performance event, so that you can analyze the performance issue and take corrective action for resolving it.

The expected range sets the performance baseline for the workload. Over time, Unified Manager learns from past performance measurements to forecast the expected performance and activity levels for the workload. The upper boundary of the expected range establishes the performance threshold. Unified Manager uses the baseline to determine when the actual latency or operations are above or below a threshold, or outside the bounds of their expected range. The comparison between the actual values and the expected values creates a performance profile for the workload.

When the actual latency for a workload exceeds the performance threshold, due to contention on a cluster component, the latency is high and the workload performs more slowly than expected. The performance of other workloads that share the same cluster components might also be slower than expected.

Unified Manager analyzes the threshold crossing event and determines whether the activity is a performance event. If the high workload activity remains consistent for a long period of time, such as several hours, Unified Manager considers the activity to be normal and dynamically adjusts the expected range to form the new performance threshold.

Some workloads might have consistently low activity, where the expected range for the operations or the latency does not have a high rate of change over time. To minimize the number of event alerts, during analysis of performance events, Unified Manager triggers an event only for low-activity volumes whose operations and latencies are much higher than expected.



In this example, the latency for a volume has an expected range, in gray, of 0 milliseconds per operation (ms/op) at its lowest and 5 ms/op at its highest. If the actual latency, in blue, suddenly increases to 10 ms/op, due to an intermittent spike in network traffic or contention on a cluster component, it is then above the expected range and has exceeded the performance threshold.

When network traffic has decreased, or the cluster component is no longer in contention, the latency returns within the expected range. If the latency remains at or above 10 ms/op for a long period of time, you might need to take corrective action to resolve the event.

How Unified Manager uses workload latency to identify performance issues

The workload latency (response time) is the time it takes for a volume on a cluster to respond to I/O requests from client applications. Unified Manager uses the latency to detect and alert you to performance events.

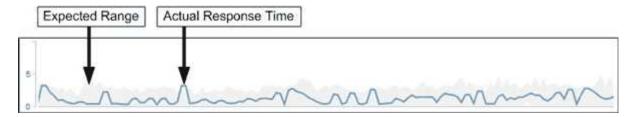
A high latency means that requests from applications to a volume on a cluster are taking longer than usual. The cause of the high latency could be on the cluster itself, due to contention on one or more cluster components. High latency could also be caused by issues outside of the cluster, such as network bottlenecks, issues with the client hosting the applications, or issues with the applications themselves.



Unified Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

Operations on the cluster, such as making backups or running deduplication, that increase their demand of cluster components shared by other workloads can also contribute to high latency. If the actual latency exceeds the performance threshold of the expected range, Unified Manager analyzes the event to determine whether it is a performance event that you might need to resolve. The latency is measured in milliseconds per operation (ms/op).

On the Performance/Volume Details page, you can view an analysis of the latency statistics to see how the activity of individual processes, such as read and write requests, compares to the overall latency statistics. The comparison helps you determine which operations have the highest activity or whether specific operations have abnormal activity that is impacting the latency for a volume. When analyzing performance events, you can use the latency statistics to determine whether an event was caused by an issue on the cluster. You can also identify the specific workload activities or cluster components that are involved in the event.



This example shows the Latency chart on the Performance/Volume Details page. The actual response time (latency) activity is a blue line and the expected range is gray.

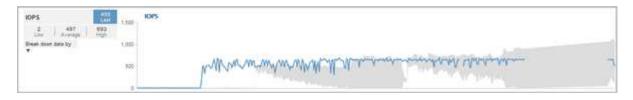


There can be gaps in the blue line if Unified Manager was unable to gather data. This can occur because the cluster or volume was unreachable, Unified Manager was turned off during that time, or the collection was taking longer than the 5 minute collection period.

How cluster operations can affect workload latency

Operations (IOPS) represent the activity of all user-defined and system-defined workloads on a cluster. The IOPS statistics help you determine whether cluster processes, such as making backups or running deduplication, are impacting workload latency (response time) or might have caused, or contributed to, a performance event.

When analyzing performance events, you can use the IOPS statistics to determine whether a performance event was caused by an issue on the cluster. You can identify the specific workload activities that might have been the main contributors to the performance event. IOPS are measured in operations per second (ops/sec).



This example shows the IOPS chart on the Performance/Volume Details page. The actual operations statistics is a blue line and the expected range of operations statistics is gray.



In some cases where a cluster is overloaded, Unified Manager might display the message Data collection is taking too long on Cluster cluster_name. This means that not enough statistics have been collected for Unified Manager to analyze. You need to reduce the resources the cluster is using so that statistics can be collected.

Performance monitoring of MetroCluster configurations

Unified Manager enables you to monitor the write throughput between clusters in a MetroCluster configuration to identify workloads with a high amount of write throughput. If these high-performing workloads are causing other volumes on the local cluster to have high I/O response times, Unified Manager triggers performance events to notify you.

When a local cluster in a MetroCluster configuration mirrors its data to its partner cluster, the data is written to NVRAM and then transferred over the interswitch links (ISLs) to the remote aggregates. Unified Manager analyzes the NVRAM to identify the workloads whose high write throughput is overutilizing the NVRAM, placing the NVRAM in contention.

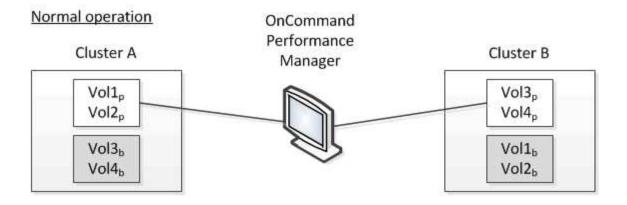
Workloads whose deviation in response time has exceeded the performance threshold are called *victims* and workloads whose deviation in write throughput to the NVRAM is higher than usual, causing the contention, are called *bullies*. Because only the write requests are mirrored to the partner cluster, Unified Manager does not analyze read throughput.

Unified Manager treats the clusters in a MetroCluster configuration as individual clusters. It does not distinguish between clusters that are partners or correlate the write throughput from each cluster.

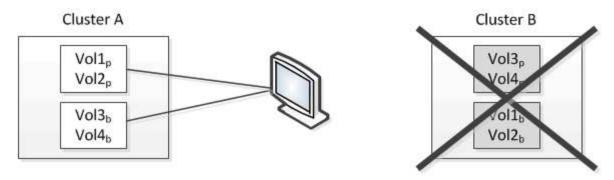
Volume behavior during switchover and switchback

Events that trigger a switchover or switchback cause active volumes to be moved from one cluster to the other cluster in the disaster recovery group. The volumes on the cluster that were active and serving data to clients are stopped, and the volumes on the other cluster are activated and start serving data. Unified Manager monitors only those volumes that are active and running.

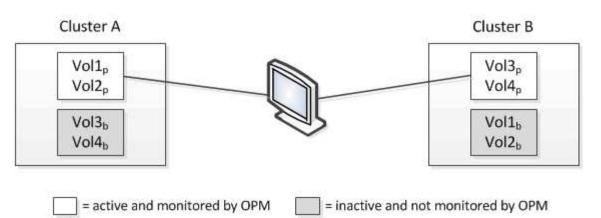
Because volumes are moved from one cluster to another, it is recommended that you monitor both clusters. A single instance of Unified Manager can monitor both clusters in a MetroCluster configuration, but sometimes the distance between the two locations necessitates using two Unified Manager instances to monitor both clusters. The following figure shows a single instance of Unified Manager:



Cluster B fails --- switchover to Cluster A



Cluster B is repaired --- switchback to Cluster B



The volumes with p in their names indicate the primary volumes, and the volumes with b in their names are mirrored backup volumes that are created by SnapMirror.

During normal operation:

- Cluster A has two active volumes: Vol1p and Vol2p.
- Cluster B has two active volumes: Vol3p and Vol4p.
- Cluster A has two inactive volumes: Vol3b and Vol4b.
- Cluster B has two inactive volumes: Vol1b and Vol2b.

Information pertaining to each of the active volumes (statistics, events, and so on) is collected by Unified Manager. Vol1p and Vol2p statistics are collected by Cluster A, and Vol3p and Vol4p statistics are collected by Cluster B.

After a catastrophic failure causes a switchover of active volumes from Cluster B to Cluster A:

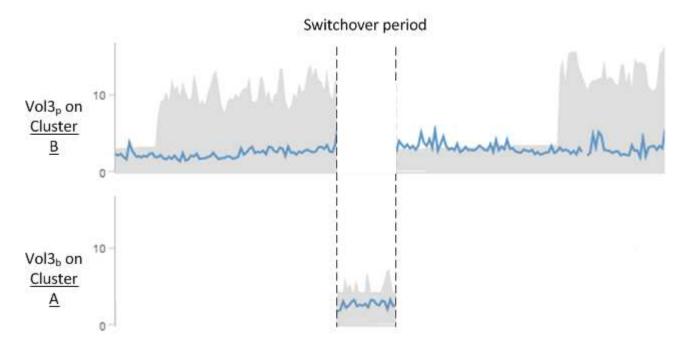
- Cluster A has four active volumes: Vol1p, Vol2p, Vol3b, and Vol4b.
- Cluster B has four inactive volumes: Vol3p, Vol4p, Vol1b, and Vol2b.

As during normal operation, information pertaining to each of the active volumes is collected by Unified Manager. But in this case, Vol1p and Vol2p statistics are collected by Cluster A, and Vol3b and Vol4b statistics are also collected by Cluster A.

Note that Vol3p and Vol3b are not the same volumes, because they are on different clusters. The information in Unified Manager for Vol3p is not the same as Vol3b:

- During switchover to Cluster A, Vol3p statistics and events are not visible.
- On the very first switchover, Vol3b looks like a new volume with no historical information.

When Cluster B is repaired and a switchback is performed, Vol3p is active again on Cluster B, with the historical statistics and a gap of statistics for the period during the switchover. Vol3b is not viewable from Cluster A until another switchover occurs:





- MetroCluster volumes that are inactive, for example, Vol3b on Cluster A after switchback, are identified with the message "This volume was deleted". The volume is not actually deleted, but it is not currently being monitored by Unified Manager because it is not the active volume.
- If a single Unified Manager is monitoring both clusters in a MetroCluster configuration, volume search returns information for whichever volume is active at that time. For example, a search for "Vol3" would return statistics and events for Vol3b on Cluster A if a switchover has occurred and Vol3 has become active on Cluster A.

Performance event analysis and notification

Performance events notify you about I/O performance issues on a volume workload caused by contention on a cluster component. Unified Manager analyzes the event to

identify all workloads involved, the component in contention, and whether the event is still an issue that you might need to resolve.

Unified Manager monitors the I/O latency (response time) and IOPS (operations) for volumes on a cluster. When other workloads overuse a cluster component, for example, the component is in contention and cannot perform at an optimal level to meet workload demands. The performance of other workloads that are using the same component might be impacted, causing their latencies to increase. If the latency crosses the performance threshold, Unified Manager triggers a performance event and sends an email alert to notify you.

Event analysis

Unified Manager performs the following analyses, using the previous 15 days of performance statistics, to identify the victim workloads, bully workloads, and the cluster component involved in an event:

- Identifies victim workloads whose latency has crossed the performance threshold, which is the upper boundary of the expected range:
 - For volumes on HDD or Flash Pool (hybrid) aggregates, events are triggered only when the latency is greater than 5 milliseconds (ms) and the IOPS are more than 10 operations per second (ops/sec).
 - For volumes on all-SSD aggregates or FabricPool (composite) aggregates, events are triggered only when the latency is greater than 1 ms and the IOPS are more than 100 ops/sec.
- Identifies the cluster component in contention.



If the latency of victim workloads at the cluster interconnect is greater than 1 ms, Unified Manager treats this as significant and triggers an event for the cluster interconnect.

- Identifies the bully workloads that are overusing the cluster component and causing it to be in contention.
- Ranks the workloads involved, based on their deviation in utilization or activity of a cluster component, to
 determine which bullies have the highest change in usage of the cluster component and which victims are
 the most impacted.

An event might occur for only a brief moment and then correct itself after the component it is using is no longer in contention. A continuous event is one that reoccurs for the same cluster component within a five-minute interval and remains in the active state. For continuous events, Unified Manager triggers an alert after detecting the same event during two consecutive analysis intervals. Events that remain unresolved, which have a state of new, can display different description messages as workloads involved in the event change.

When an event is resolved, it remains available in Unified Manager as part of the record of past performance issues for a volume. Each event has a unique ID that identifies the event type and the volumes, cluster, and cluster components involved.



A single volume can be involved in more than one event at the same time.

Event state

Events can be in one of the following states:

Active

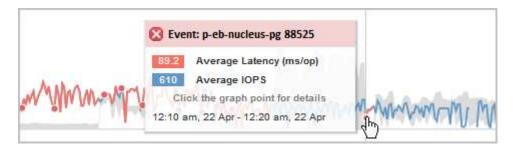
Indicates that the performance event is currently active (new or acknowledged). The issue causing the event has not corrected itself or has not been resolved. The performance counter for the storage object remains above the performance threshold.

Obsolete

Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the storage object is no longer above the performance threshold.

Event notification

The event alerts are displayed on the Dashboards/Overview page, Dashboards/Performance page, Performance/Volume Details page, and they are sent to specified email addresses. You can view detailed analysis information about an event and get suggestions for resolving it on the Event details page.



In this example, an event is indicated by a red dot () on the Latency chart on the Performance/Volume Details page. Hovering your mouse cursor over the red dot displays a popup with more details about the event and options for analyzing it.

Event interaction

On the Performance/Volume Details page, you can interact with events in the following ways:

• Moving the pointer over a red dot displays a message that shows the event ID, along with the latency, number of operations per second, and the date and time when the event was detected.

If there are multiple events for the same time period, the message shows the number of events, along with the average latency and operations per second for the volume.

 Clicking a single event displays a dialog box that shows more detailed information about the event, including the cluster components that are involved, similar to the Summary section on the Event details page.

The component in contention is circled and highlighted red. You can click either the event ID or **View full analysis** to view the full analysis on the Event details page. If there are multiple events for the same time period, the dialog box shows details about the three most recent events. You can click an event ID to view the event analysis on the Event details page. If there are more than three events for the same time period, clicking the red dot does not display the dialog box.

How Unified Manager determines the performance impact for an event

Unified Manager uses the deviation in activity, utilization, write throughput, cluster component usage, or I/O latency (response time) for a workload to determine the level of impact to workload performance. This information determines the role of each workload in the event and how they are ranked on the Event details page.

Unified Manager compares the last analyzed values for a workload to the expected range of values. The difference between the values last analyzed and the expected range of values identifies the workloads whose

performance was most impacted by the event.

For example, suppose a cluster contains two workloads: Workload A and Workload B. The expected range for Workload A is 5-10 milliseconds per operation (ms/op) and its actual latency is usually around 7 ms/op. The expected range for Workload B is 10-20 ms/op and its actual latency is usually around 15 ms/op. Both workloads are well within their expected range for latency. Due to contention on the cluster, the latency of both workloads increases to 40 ms/op, crossing the performance threshold, which is the upper bounds of the expected range, and triggering events. The deviation in latency, from the expected values to the values above the performance threshold, for Workload A is around 33 ms/op, and the deviation for Workload B is around 25 ms/op. The latency of both workloads spike to 40 ms/op, but Workload A had the bigger performance impact because it had the higher latency deviation at 33 ms/op.

On the Event details page, in the System Diagnosis section, you can sort workloads by their deviation in activity, utilization, or throughput for a cluster component. You can also sort workloads by latency. When you select a sort option, Unified Manager analyzes the deviation in activity, utilization, throughput, or latency since the event was detected from the expected values to determine the workload sort order. For the latency, the red dots () indicate a performance threshold crossing by a victim workload, and the subsequent impact to the latency. Each red dot indicates a higher level of deviation in latency, which helps you identify the victim workloads whose latency was impacted the most by an event.

Cluster components and why they can be in contention

You can identify cluster performance issues when a cluster component goes into contention. The performance of volume workloads that use the component slow down and their response time (latency) for client requests increases, which triggers an event in Unified Manager.

A component that is in contention cannot perform at an optimal level. Its performance has declined, and the performance of other cluster components and workloads, called *victims*, might have increased latency. To bring a component out of contention, you must reduce its workload or increase its ability to handle more work, so that the performance can return to normal levels. Because Unified Manager collects and analyzes workload performance in five-minute intervals, it detects only when a cluster component is consistently overused. Transient spikes of overusage that last for only a short duration within the five-minute interval are not detected.

For example, a storage aggregate might be under contention because one or more workloads on it are competing for their I/O requests to be fulfilled. Other workloads on the aggregate can be impacted, causing their performance to decrease. To reduce the amount of activity on the aggregate, there are different steps you can take, such as moving one or more workloads to a less busy aggregate, to lessen the overall workload demand on the current aggregate. For a QoS policy group, you can adjust the throughput limit, or move workloads to a different policy group, so that the workloads are no longer being throttled.

Unified Manager monitors the following cluster components to alert you when they are in contention:

Network

Represents the wait time of I/O requests by the iSCSI protocols or the Fibre Channel (FC) protocols on the cluster. The wait time is time spent waiting for iSCSI Ready to Transfer (R2T) or FCP Transfer Ready (XFER_RDY) transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the block protocol layer is impacting the latency of one or more workloads.

Network Processing

Represents the software component in the cluster involved with I/O processing between the protocol layer

and the cluster. The node handling network processing might have changed since the event was detected. If the network processing component is in contention, it means high utilization at the network processing node is impacting the latency of one or more workloads.

QoS Policy

Represents the storage Quality of Service (QoS) policy group of which the workload is a member. If the policy group component is in contention, it means all workloads in the policy group are being throttled by the set throughput limit, which is impacting the latency of one or more of those workloads.

Cluster Interconnect

Represents the cables and adapters with which clustered nodes are physically connected. If the cluster interconnect component is in contention, it means high wait time for I/O requests at the cluster interconnect is impacting the latency of one or more workloads.

Data Processing

Represents the software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains the workload. The node handling data processing might have changed since the event was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the latency of one or more workloads.

MetroCluster Resources

Represents the MetroCluster resources, including NVRAM and interswitch links (ISLs), used to mirror data between clusters in a MetroCluster configuration. If the MetroCluster component is in contention, it means high write throughput from workloads on the local cluster or a link health issue is impacting the latency of one or more workloads on the local cluster. If the cluster is not in a MetroCluster configuration, this icon is not displayed.

Aggregate or SSD Aggregate Ops

Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate). An "SSD Aggregate" consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

Cloud Latency

Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

Sync SnapMirror

Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

Roles of workloads involved in a performance event

Unified Manager uses roles to identify the involvement of a workload in a performance event. The roles include victims, bullies, and sharks. A user-defined workload can be a victim, bully, and shark at the same time.

Role	Description
Victim	A user-defined workload whose performance has decreased due to other workloads, called bullies, that are over-using a cluster component. Only user-defined workloads are identified as victims. Unified Manager identifies victim workloads based on their deviation in latency, where the actual latency, during an event, has greatly increased from its expected range of latency.
Bully	A user-defined or system-defined workload whose over-use of a cluster component has caused the performance of other workloads, called victims, to decrease. Unified Manager identifies bully workloads based on their deviation in usage of a cluster component, where the actual usage, during an event, has greatly increased from its expected range of usage.
Shark	A user-defined workload with the highest usage of a cluster component compared to all workloads involved in an event. Unified Manager identifies shark workloads based on their usage of a cluster component during an event.

Workloads on a cluster can share many of the cluster components, such as storage aggregates and the CPU for network and data processing. When a workload, such as a volume, increases its usage of a cluster component to the point that the component cannot efficiently meet workload demands, the component is in contention. The workload that is over-using a cluster component is a bully. The other workloads that share those components, and whose performance is impacted by the bully, are the victims. Activity from system-defined workloads, such as deduplication or Snapshot copies, can also escalate into "bullying".

When Unified Manager detects an event, it identifies all workloads and cluster components involved, including the bully workloads that caused the event, the cluster component that is in contention, and the victim workloads whose performance has decreased due to the increased activity of bully workloads.



If Unified Manager cannot identify the bully workloads, it only alerts on the victim workloads and the cluster component involved.

Unified Manager can identify workloads that are victims of bully workloads, and also identify when those same workloads become bully workloads. A workload can be a bully to itself. For example, a high-performing workload that is being throttled by a policy group limit causes all workloads in the policy group to be throttled, including itself. A workload that is a bully or a victim in an ongoing performance event might change its role or no longer be a participant in the event. On the Performance/Volume Details page, in the Events List table, when the selected volume changes its participant role, the date and time of the role change is displayed.

Analyzing workload performance

Unified Manager enables you to monitor and analyze I/O performance of volume workloads on your clusters. You can determine whether a performance issue is on the cluster and whether storage is the issue.



This chapter describes how to analyze workload performance using the Performance/Volume Details page and the Event details page.

Determining whether a workload has a performance issue

You can use Unified Manager to determine whether a detected performance event was truly caused by a performance issue on the cluster. The event might have been caused a spike in activity, for example, that drove up its response time, but now the response time has returned to it usual levels.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have identified the name of the volume, or associated LUN, you want to analyze.
- Unified Manager must have collected and analyzed a minimum of five days of performance statistics from the cluster.

About this task

If you are viewing the Event details page, you can click the name link for a volume to go directly to the Performance/Volume Details page.

Steps

1. In the **Search** bar, type at least the first three characters of the volume name.

The name of the volume is displayed in the search results.

2. Click the name of the volume.

The volume is displayed on the Performance/Volume Details page.

- 3. In the **Historic data** chart, click **5d** to display the last five days of historical data.
- 4. Review the **Latency** chart to answer the following questions:
 - Are there new performance events?
 - Are there historic performance events, indicating that the volume has had issues in the past?
 - · Are there spikes in the response time, even if the spikes are within the expected range?
 - On the Have there been configuration changes on the cluster that might have impacted performance? If the response time for the volume does not display performance events, spikes in activity, or recent configuration changes that might have impacted the response time, you can rule out the performance issue being caused by the cluster.

Investigating a perceived slow response time for a workload

You can use Unified Manager to determine whether operations on the cluster might have contributed to the slow response time (latency) for a volume workload.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have identified the name of the volume, or associated LUN, you want to analyze.
- Unified Manager must have collected and analyzed a minimum of five days of performance statistics from the cluster.

About this task

If you are viewing the Event details page, you can click the name for a volume to go directly to the Performance/Volume Details page.

Steps

1. In the **Search** bar, type the name of the volume.

The name of the volume is displayed in the search results.

2. Click the name of the volume.

The volume is displayed on the Performance/Volume Details page.

- 3. On the Historic data chart, click 5d to display the last five days of historical data.
- 4. Review the **IOPS** chart to answer the following questions:
 - Are there dramatic spikes in the activity?
 - · Are there dramatic drops in the activity?
 - Are there abnormal changes in the operations pattern? If the operations do not display dramatic spikes or drops in activity, and there were no changes to the cluster configuration during this time, the storage administrator can confirm that other workloads have not impacted volume performance.
- 5. On the Break down data by menu, under IOPS, select Reads/writes/other.
- 6. Click Submit.

The Reads/writes/other chart is displayed below the IOPS chart.

7. Review the **Reads/writes/other** chart to identify dramatic spikes or drops in the amount of reads or writes for the volume.

If there are no dramatic spikes or drops in reads or writes, the storage administrator can confirm that I/O on the cluster is operating normally. Any performance issues might be on the network or the connected clients.

Identifying trends of I/O response time on cluster components

You can use Unified Manager to view the performance trends for all monitored cluster

components for a volume workload. You can see, over time, which components have the highest usage, whether the highest usage is from read or write requests, and how the usage has impacted the workload response time.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have identified the name of the volume or associated LUN you want to analyze.
- To display 30 days of performance statistics, Unified Manager must have collected and analyzed a minimum of 30 days of performance statistics from the cluster.

About this task

Identifying performance trends for the cluster components helps the administrator decide whether the cluster is being overused or underused.

If you are viewing the Event details page, you can click the name for a volume to go directly to the Performance/Volume Details page.

Steps

1. In the **Search** bar, type the name of the volume.

The name of the volume is displayed in the search results.

2. Click the name of the volume.

The volume is displayed on the Performance/Volume Details page.

- 3. On the Historic data chart, click **30d** to display the last 30 days of historical data.
- 4. Click Break down data by.
- 5. Under Latency, select Cluster Components and Reads/writes latency.
- 6. Click Submit.

Both charts are displayed below the Latency chart.

7. Review the **Cluster Components** chart.

The chart breaks down the total response time by cluster component. The response time at the aggregate is the highest.

8. Compare the **Cluster Components** chart to the **Latency** chart.

The Latency chart shows spikes in the total response time that are aligned with the spikes in response time for the aggregate. There are a few at the end of the 30-day time frame, where the performance threshold was crossed.

Review the Reads/writes latency chart.

The chart shows a higher response time for write requests than read requests, indicating that the client applications are waiting longer than usual to have their write requests fulfilled.

10. Compare the Reads/writes latency chart to the Latency chart.

The spikes in total response time that align with the aggregate in the Cluster Components chart also align with the writes in the Reads/writes latency chart. The administrator must decide whether the client applications using the workload must be addressed or whether the aggregate is being overused.

Analyzing the performance improvements achieved from moving a volume

You can use Unified Manager to investigate the impact of a volume move operation on the latency (response time) of other volumes on the cluster. Moving a high performing volume to a less busy aggregate or an aggregate with flash storage enabled allows the volume to perform more efficiently.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have identified the name of the volume, or associated LUN, you want to analyze.
- Unified Manager must have collected and analyzed seven days of data.

About this task

Unified Manager identifies when a volume moves between aggregates. It can detect when the volume move is occurring, completed, or failed. The Performance/Volume Details page displays a change event icon for each state of the volume move, which helps you track when a move operation occurred and helps you determine whether it might have contributed to a performance event.

If you are viewing the Event details page, you can click the name of a volume to go directly to the Performance/Volume Details page.

Steps

- 1. In the **Search** bar, type the name of the volume.
- 2. Click the name of the volume.

The volume is displayed on the Performance/Volume Details page.

- 3. In the **Historic data** chart, adjust the sliders to show activity from the previous work week.
- 4. Analyze the **Latency** chart and the **IOPS** chart to see how the volume performed over the last few days.

Assume that you notice a consistent pattern of very high average response times of over 42 milliseconds per operation (ms/op), with performance events, each day of the week and decide to move the volume to a less busy aggregate to improve performance. Using OnCommand System Manager, you can move the volume to an aggregate with Flash Pool enabled for an increased performance boost. Approximately an hour after the volume move has been completed, you can return to Unified Manager to confirm that the move operation was completed successfully and that the latency has improved.

- 5. If the Performance/Volume Details page is not displayed, search for the volume you want to view.
- 6. On the **Historic data** chart, click **1d** to view the activity from the last one day, a few hours since the volume move was completed.

At the bottom of the page, in the Events time line, a change event icon () is displayed to indicate the

time that the volume move operation was completed. A black, vertical line is also displayed from the change event icon to the Latency chart.

7. Point your cursor to the change event icon to view details about the event in the Events List.

Because the volume moved to an aggregate with Flash Pool enabled, you can see the change in read and write I/O to cache.

8. On the Break down data by menu, under MBps, select Cache hit ratio.

The Cache hit ratio chart displays statistics about the reads and writes to cache.

The volume successfully moved to a less busy aggregate and the change event is highlighted in the Events List on the right. The average latency decreased significantly from over 42 ms/op to around 24 ms/op. The current latency is around 1.5 ms/op. In the Cache hit ratio chart, the amount of successful read and write hits to cache is now at 100% because the volume is now on an aggregate with Flash Pool enabled.

How moving a FlexVol volume works

Knowing how moving a FlexVol volume works helps you to determine whether the volume move satisfies service-level agreements and to understand where a volume move is in the volume move process.

FlexVol volumes are moved from one aggregate or node to another within the same storage virtual machine (SVM). A volume move does not disrupt client access during the move.

Moving a volume occurs in multiple phases:

- · A new volume is made on the destination aggregate.
- The data from the original volume is copied to the new volume.

During this time, the original volume is intact and available for clients to access.

• At the end of the move process, client access is temporarily blocked.

During this time the system performs a final replication from the source volume to the destination volume, swaps the identities of the source and destination volumes, and changes the destination volume to the source volume.

 After completing the move, the system routes client traffic to the new source volume and resumes client access.

The move is not disruptive to client access because the time in which client access is blocked ends before clients notice a disruption and time out. Client access is blocked for 35 seconds by default. If the volume move operation cannot finish in the time that access is denied, the system aborts this final phase of the volume move operation and allows client access. The system attempts the final phase three times by default. After the third attempt, the system waits an hour before attempting the final phase sequence again. The system runs the final phase of the volume move operation until the volume move is complete.

Performance/Volume Details page

This page provides detailed performance statistics for all I/O activity and operations for

the selected FlexVol volume, FlexGroup volume, or FlexGroup constituent workload. You can select a specific time frame over which to view the statistics and events for the volume. The events identify performance events and changes that might be impacting I/O performance.

Historic data chart

Plots the historical performance analysis data for the selected volume. You can click and drag the sliders to specify a time frame. The sliders increase and decrease the time frame window. The data outside the time frame window is grayed out. You can use the slider at the bottom of the chart to move the time frame window across the historical data. The entire page, including the displayed charts and events, reflects the data available within the time frame window. Unified Manager retains a maximum of 30 days of historical data on this page.



On the historic data chart, if you select a time frame of more than 1 day, depending on your screen resolution, the charts display the maximum values for response time and IOPS across the number of days.

Options

· Time selector

Specifies the time range over which to view the volume performance statistics for the entire page. You can click 1 day (1d) through 30 days (30d), or click **Custom** to select a custom range. For a custom range, you can select a beginning and end date, and then click **Update** to update the entire page.



If you access the Performance/Volume Details page by clicking the name link of a volume on the Event details page, a time range, such as 1 day or 5 days prior to the current day, is automatically selected by default. When you move the slider in the historic data chart, the time range changes to a custom range, but the **Custom** time selector is not selected. The default time selector remains selected.

Break down data by

Provides a list of charts you can add to the Performance/Volume Details page to display more detailed performance statistics for the selected volume.

Performance statistics displayed in the data breakdown charts

You can use the graphs to view performance trending for a volume. You can also view statistics for reads and writes, network protocol activity, the impact of QoS policy group throttling on latency, the ratio of reads and writes to cache storage, the total cluster CPU time used by a workload, and specific cluster components.

These views display a maximum of 30 days of statistics from the current day. On the historic data chart, if you select a time frame of more than 1 day, depending on your screen resolution, the charts display the maximum values for latency and IOPS across the number of days.



You can use the Select All check box to select, or deselect, all the listed chart options.

Latency

The following charts detail the latency, or response time, information for the selected workload:

Cluster Components

Displays a graph of the time spent at each cluster component used by the selected volume.

The chart helps you determine the latency impact by each component as it relates to the total latency. You can use the check box next to each component to show and hide its graph.

For QoS policy groups, data is only displayed for user-defined policy groups. Zeros are displayed for system-defined policy groups, such as default policy groups.

Reads/writes latency

Displays a graph of the latencies of the successful read and write requests from the selected volume workload over the selected time frame.

Write requests are an orange line and read requests are a blue line. The requests are specific to the latency for the selected volume workload, not all workloads on the cluster.



The read and write statistics might not always add up to the total latency statistics displayed in the Latency chart. This is expected behavior based on how Unified Manager collects and analyzes read and write statistics for a workload.

Policy Group Impact

Displays a graph of the percentage of the latency for the selected volume workload that is impacted by the throughput limit on its QoS policy group.

If the workload is throttled, the percentage indicates how much the throttling contributed to the latency at a specific point in time. The percentage values indicate the amount of throttling:

- 0% = no throttling
- > 0% = throttling
- > 20% = critical throttling
 If the cluster can handle more work, you can reduce throttling by increasing the policy group limit.
 Another option is to move the workload to a less busy aggregate.



The chart displays for workloads in a user-defined QoS policy group with a set throughput limit only. It does not display if the workloads are in a system-defined policy group, such as the default policy group, or a policy group that does not have a QoS limit. For a QoS policy group, you can point the cursor to the name of the policy group to display its throughput limit and the last time it was modified. If the policy group was modified before the associated cluster was added to Unified Manager, the last modified time is the date and time when Unified Manager first discovered the cluster.

· IOPS

The following charts detail the IOPS data for the selected workload:

Reads/writes/other

Displays a graph showing the number of read and write IOPS and other IOPS, per second, over the selected time frame.

Other IOPS are protocol activities initiated by the client that are not reads or writes. For example, in an NFS environment, this could be metadata operations such as getattr, setattr, or fsstat. In a CIFS environment, this could be attribute lookups, directory listings, or antivirus scans. Write IOPS are an orange line and read requests are a blue line. The requests are specific to all operations for the selected volume workload, not all operations on the cluster.

MBps

The following charts detail the throughput data for the selected workload:

Cache hit ratio

Displays a graph of the percentage of read requests from client applications satisfied by cache over the selected time frame.

The cache could be on Flash Cache cards or solid state drives (SSDs) in Flash Pool aggregates. A cache hit, in blue, is a read from cache. A cache miss, in orange, is a read from a disk in the aggregate. The requests are specific to the selected volume workload, not all workloads on the cluster.

You can view more detailed information about volume cache usage in the Unified Manager Health pages and in OnCommand System Manager.

Components

The following charts detail the data by cluster component used by the selected workload:

Cluster CPU Time

Displays a graph of the CPU usage time, in ms, for all nodes in the cluster used by the selected workload.

The graph displays the combined CPU usage time for network processing and data processing. The CPU time for system-defined workloads that are associated to the selected workload, and are using the same nodes for data processing, is also included. You can use the chart to determine whether the workload is a high consumer of the CPU resources on the cluster. You can also use the chart, in combination with the Reads/writes latency chart under the Latency chart, or the Reads/writes/other chart under the IOPS chart, to determine how changes to workload activity over time impact cluster CPU utilization.

Disk Utilization

Displays a graph showing the percentage of utilization on the data disks in the storage aggregate over the selected time frame.

The utilization includes disk read and write requests from the selected volume workload only. Reads from cache are not included. The utilization is specific to the selected volume workload, not all workloads on the disks. If a monitored volume is involved in a volume move, the utilization values in this chart are for the target aggregate to which the volume moved.

How graphs of performance data work

Unified Manager uses graphs or charts to show you volume performance statistics and events over a specified period of time.

The graphs enable you to customize the range of time for which to view data. The data is displayed with the time frame on the horizontal axis of the graph and the counters on the vertical axis, with point intervals along the graph lines. The vertical axis is dynamic; the values adjust based on the peaks of the expected or actual values.

Selecting time frames

On the Performance/Volume Details page, the Historic data chart enables you to select a time frame for all graphs on the page. The 1d, 5d, 10d, and 30d buttons specify 1 day through 30 days (1 month) and the **Custom** button enables you to specify a custom time range within that 30 days. Each point on a graph represents a 5-minute collection interval, and a maximum of 30 days of historical performance data is retained. Note that intervals also account for network delays and other anomalies.



In this example, the Historic data chart has a time frame set to the beginning and the end of the month of March. In the selected time frame, all historic data before March is grayed out.

Viewing data point information

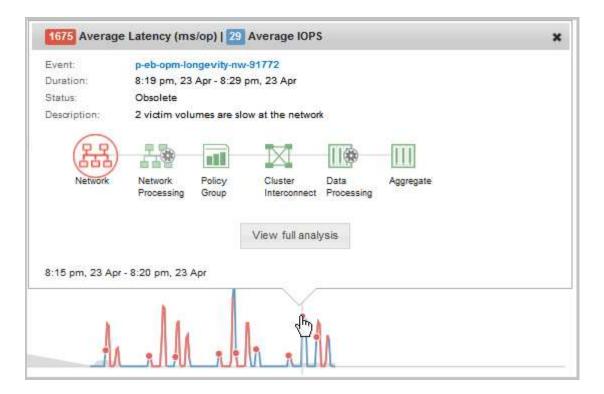
To view data point information on a graph, you can position the cursor over a specific point within the graph, and a pop-up box displays listing the value and date and time information.



In this example, positioning the cursor over the IOPS chart on the Performance/Volume Details page displays the response time and operations values between 3:50 a.m. and 3:55 a.m. on October 20th.

Viewing performance event information

To view event information on a graph, you can position the cursor over an event icon to view summary information in a pop-up box, or you can click the event icon for more detailed information.



In this example, on the Performance/Volume Details page, clicking an event icon on the Latency chart displays detailed information about the event in a pop-up box. The event is also highlighted in the Events List.

Managing reports

OnCommand Unified Manager enables you to create and manage reports so that you can view customized information about the capacity, utilization, and performance of storage objects and events related to storage objects.

The implementation of scheduling and generating reports from the Reports option in the left Navigation pane will be deprecated in a future release. You can extract health and performance data from Unified Manager using these additional methods:



- Extracting data directly from the grid pages in the GUI
- · Using Open Database for access to all the available objects
- Executing Unified Manager REST APIs to return the information you are interested in reviewing

Scheduling reports

You can schedule your reports from the Report details page and email the scheduled reports to one or more recipients in a particular format at a specified frequency. For example, you can schedule a report to be sent as email, in the PDF format, every Monday.

Steps

1. In the left navigation pane, click **Reports**, and then click **Run Report** for the report that you want.

The Report details page is displayed.

- 2. Click Actions > Schedule Report.
- 3. In the Schedule Report dialog box, select one of the preferred schedules for your report:

If you want to	Then
Select a schedule from the existing list of schedules	Click Use Existing Schedule and select the schedule.
Create a new schedule	Click Create New Schedule , and then enter the schedule name, specify the email address, select the report format and frequency, and enter a time or day you want the report to run. You can specify one or more email addresses, separated by commas. The <i>PDF</i> option is selected as the default report format. The <i>Hourly</i> option is selected as the default frequency.

4. Click Schedule.

Sharing reports

You can email and share your reports with one or more users.

Steps

1. In the left navigation pane, click Reports, and then click Run Report for the report that you want.

The Report details page is displayed.

- 2. Click Actions > Share.
- In the Share Report dialog box, specify the email address of the recipient with whom you want to share the report.

You can specify one or more email addresses, separated by commas.

4. Specify the subject of the email.

By default, the name of the report appears as the subject of the email.

5. Select the report format.

The *PDF* option is selected as the default report format. If the XHTML format is selected, the recipient must open the report that is sent by email by using a supported web browser.

6. Click Share.

Managing report schedules

You can manage your report schedules from the Manage Report Schedules dialog box. You can add a new schedule and view, modify, or delete existing schedules.

Steps

- 1. In the left navigation pane, click **Reports**, and then click **Manage Report Schedules**.
- 2. In the Manage Report Schedules dialog box,

If you want to	Then
View or modify an existing schedule	 a. Select the schedule from the list displayed in the left pane. The schedule details are displayed. b. Make the necessary changes. c. Click Save or Save and Close.
Delete an existing schedule	a. Select the schedule from the list displayed in the left pane.The schedule details are displayed.b. Click Delete Schedule.
Add a new schedule	 a. Click Add Schedule. b. A new schedule form appears in the right pane. c. Enter the schedule name, recipient email address, report format and frequency and select the reports you want to schedule. d. Click Save. The new schedule will be added in the Schedules list.

Customizing a report

You can customize reports in the Report details page and then save the customized report with a different name.

About this task

After you save a customized report you cannot modify any of the filters you applied to create the report because the report is considered "new". So make sure you are satisfied with all changes before you save the report. You can, however, apply new filters to the report.

- 1. In the left navigation pane, click **Reports**, and then click **Run Report** for the report that you want to customize.
 - The Report details page is displayed.
- 2. Customize the report as necessary, and then click **Actions** > **Save Customized Report As**.

3. In the **Save Customized Report As** dialog box, enter a name for the customized report and a brief description about the customization so that others will understand what the report displays.

By default, the current report name is displayed.

4. Click Save.

If you receive the error message "Failed to save the custom report. The required file was not created", wait a few moments, and then click **Save** again. This issue has been seen when there is a slow connection between the web browser and the Unified Manager server.

Results

The customized report is saved and displayed in its respective report category in the Report details page.

Editing a customized report

You can make additional changes to an already customized report and save the report. You cannot change the name of the report after you have saved it.

Steps

- 1. In the left navigation pane, click **Reports**, and then click **Run Report** for the report that you want to edit.
 - The Report details page is displayed.
- 2. Modify the report as necessary, and then click **Actions** > **Save Custom Report**.
- In the Save Custom Report dialog box, enter a brief description about the changes made on the custom report and click Save.

Importing reports

If you have created a report outside of Unified Manager, you can import and save the report file to use with Unified Manager.

Before you begin

You must have the OnCommand Administrator role.

You must ensure that the report you plan to import is supported by Unified Manager.

Steps

- 1. In the left navigation pane, click **Reports**, and then click **Import Report**.
- 2. In the **Import Report** dialog box, click **Browse** and select the file you want to import, and then enter a name and brief description of the report.
- 3. Click **Import**.

If you cannot import the report, you can check the log file to find the error causing the issue.

Understanding more about reports

You can use the option to run, delete, export, and import reports. You can also create custom reports and save the customized report. You can perform additional operations such as filtering, sorting, grouping, and formatting.

What reports do

Reports display detailed information about storage objects, which enable you to review and identify potential issues.

You can save, delete, share, schedule, and import reports. You can also search for specific reports. You can customize reports to address specific use cases, and save the customized report for future use. You can perform additional operations such as filtering, sorting, grouping, and formatting.

By default, each report group is displayed by report type and description. You can run reports to view a specific report group.

After you run a report, you can further customize it and save the customized report. You can view the custom reports that are saved in the Reports page, grouped under the specific report category.

You can schedule reports to be sent, or share reports in one of the supported formats: PDF, XHTML, CSV, XLS, or text.

You can export reports in different formats and save them on your desktop. You can export individual column data from the generated reports.

You can import report design files (.rptdesign files), and save the imported reports in the Reports page. You can delete custom and imported reports.

You can import the following reports:

- Reports with multiple headers that have a column span set to one
- Reports with charts only
- Reports with lists and grid only

Reports in text, CSV, and Excel formats are supported in the following scenarios:

- Table element only in the .rptdesign file
- · A table with just one header as a row

You cannot import reports that have a column span of more than one. If a report in text, CSV, or Excel format has more than a one-header row, only the first header row is considered, and the remaining rows are ignored.

Unified Manager databases accessible for custom reporting

Unified Manager uses a MySQL database to store data from the clusters that it is monitoring. Data is persisted into various schemas in the MySQL database.

Starting with Unified Manager 7.3, additional schemas are exposed that provide access to additional table data.

All table data from the following databases are available:

Database	Description
netapp_model_view	Data about the objects on ONTAP controllers.
netapp_performance	Cluster specific performance counters.
ocum	Unified Manager application data and information to support UI filtering, sorting, and the calculation of some derived fields.
ocum_report	Data for inventory configuration and capacity-related information.
ocum_report_birt	Same as above, but this database is consumed by built-in BIRT reports.
opm	Performance configuration settings and threshold information.
scalemonitor	Data about the Unified Manager application health and performance issues.

A reporting user — a Database user with the Report Schema role — is able to access the data in these tables. This user has read-only access to reporting and other database views directly from the Unified Manager database. Note that this user does not have permission to access any tables that contain user data or cluster credential information.

See the Technical Report for Unified Manager Reporting (TR-4565) for more details.

What report scheduling is

You can schedule a report to be generated at a specific date and time by using the **Schedule** option. The report is automatically sent by email to one or more recipients as per the schedule.

By scheduling a report, you can minimize the effort of generating and sending the reports manually. You can ensure that the current status of the storage is monitored at specified intervals by the administrators who are not otherwise notified by Unified Manager.

What report sharing is

You can share a report with one or more users through email using the **Share** option.

You must save the report prior to sharing it to ensure that the recent changes you made to the report is displayed.

You can share the report in any desired format. The **Share** option helps you to share reports through email instantly, even with persons who do not have access to Unified Manager but has a valid email address.

What report importing is

You can import a report using the **Import Report** option from Unified Manager and save the imported report with a name and a brief description. By importing reports, you can add custom reports to your environment in addition to the standard reports provided in Unified Manager.

You can import a .rptdesign file that is already created. You can run, share, schedule, and delete an imported report.

Unified Manager stores the import report log files in the files jboss.log, ocum-report.log, and ocumserver-debug.log.



Customer support will not assist with designing reports, but they will support you with issues faced during a report import operation.

The import report feature includes the following support:

- Reports with multiple headers, in which the column span is set to 1 (colspan=1)
- · Reports with charts only
- · Reports with lists and grid only
- Passwords used in reports must be encoded using "base64" format. Reports using other encoding, for example, "jce" format, will cause an error during the import process.
- Reports containing data aggregation should include the aggregated column element in the table data of the report.

Reports in text, CSV, and Excel formats are supported in the following scenarios:

- Table element only in the .rptdesign file
- A table with only one header row



You cannot import reports that have a column span of more than 1. If a report in text, CSV, or Excel format has more than one-header row, only the first header row is considered, and the rest are ignored.

Report customizations

You can customize various Unified Manager reports based on storage and utilization capacity, events, cluster inventory, NFS exports, or SVM inventory.

Storage Summary report customizations

You can customize Storage Summary reports to view and analyze information about storage capacity in HA pairs. You can use filters to display storage utilization by cluster model, capacity of the most unassigned LUNs, and capacity of available HA pairs to provision new volumes and LUNs.

Customizing the Storage Summary report to view capacity by cluster models

You can customize the Storage Summary report to analyze storage capacity and utilization of clusters, and to view aggregates included in the total raw capacity.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by cluster, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To group the report by the model name, perform the following steps:
 - a. Click in the **Model** column and click the icon.
 - b. Select Group > Add Group.
- 3. To add aggregates to the total raw capacity, perform the following steps:
 - a.
 Click in the **Total Raw Capacity** column and click the icon.
 - b. Select Aggregation.
 - c. In the **Aggregation** dialog box, clear the **table level** check box and select the **group level** check box.
 - d. Enter a label name in the **Enter Label** field, if required.
- 4. Click OK.
- 5. To add aggregates to the other columns in the report, repeat Steps 3 and 4.

Customizing the Storage Summary report to analyze cluster capacity based on the ONTAP version

You can customize the Storage Summary report to group clusters by ONTAP version, and to view aggregates relating to your total raw capacity.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

- 1. To remove grouping by cluster, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click (menu icon).
 - c. Select Group > Delete Inner Group option.

- 2. To group the report by the ONTAP version, perform the following steps:
 - a.
 Click in the **OS version** column and select the icon.
 - b. Select Group > Add Group.
- 3. To add aggregates to the total raw capacity, perform the following steps:
 - a.

 Click in the **Total Raw Capacity** column and click the icon.
 - b. Select Aggregation.
 - c. In the **Aggregation** dialog box, clear the **table level** check box and select the **group level** check box.
 - d. Enter a label name in the Enter Label field, if required.
- Click OK.

Customizing the Storage Summary report to analyze clusters with the most unallocated LUN capacity

You can customize the Storage Summary report to analyze the storage utilization of clusters, which enables you to locate the LUNs with the most unallocated capacity.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove grouping by cluster, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select **Group > Delete Inner Group**.
- 2. To sort HA pairs that have the most unallocated LUN capacity, click in the **Unallocated LUN Capacity (TB)** column, and click the icon.
- 3. Select Filter > Top/Bottom N.
- 4. In the Top/Bottom N dialog box, select Top N from the Filter field and enter a value in the text field.
- 5. Click OK.

Customizing the Storage Summary report to analyze HA pairs for available capacity to provision new volumes and LUNs

You can customize the Storage Summary report to display available HA pairs that have capacity, so that you can provision new volumes and LUNs. The report displays HA pairs sorted in order of decreasing aggregate unused capacity.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove grouping by cluster, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- To sort HA pairs with available capacity, click in the Aggregate Unused Capacity (TB) column, and click the icon.
- 3. Select Filter > Top/Bottom N.
- 4. In the Top/Bottom N dialog box, select Top N from the Filter field and enter a value in the text field.
- 5. Click OK.

Aggregate Capacity and Utilization Report customizations

You can customize reports to display a variety of information about aggregates.

Customizing the Aggregate Capacity and Utilization report to view aggregates reaching full capacity

You can customize the Aggregate Capacity and Utilization report to display aggregates sorted by increasing order of aggregate capacity utilization. This enables you to view the aggregates reaching full capacity.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove grouping by cluster and by HA pair, perform the following steps:
 - a. Click in the columns that need to be ungrouped.
 - b. Click the icon.
 - c. Select **Group > Delete Inner Group option**.
- 2. To sort the aggregates reaching full capacity, click in the **Days To Full** column, and click the icon.
- Select Filter > Top/Bottom N.
- 4. In the **Top/Bottom N** dialog box, select **Bottom N** from the **Filter** field and enter a value in the text field.
- 5. Click OK.

Customizing the Aggregate Capacity and Utilization report to display aggregates with the nearly full threshold breached

You can customize the Aggregate Capacity and Utilization report to display the top aggregates, sorted by decreasing order of Snapshot copy overflow percentage. This enables you to view the storage space still available in the aggregates.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by cluster or HA pair, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- To display the difference between the used data percentage and the nearly full threshold, add a new column:
 - a. Select a column and click the icon.
 - b. Select Column > New Computed Column.
 - c. In the New Computed Column dialog box, enter a column label.
 - d. From the Select Category list, select Math.
 - e. From the Select Function list, select DIFFERENCE.
 - f. From the Column 1 list, select Space Nearly Full Threshold (%).
 - g. From the Column 2 list, select Used Data%.
 - h. Click OK.
- 3. To filter the values greater than 0 in the new column, click in the **New computed column** and open the **Filter** dialog box by clicking the icon.
- 4. From the **Condition** drop-down list, select **Greater Than**.
- 5. In the Value field, type 0 and click OK.
- 6. To sort the values, click in the **New computed column** and click the icon.
- 7. Select Filter > Top/Bottom N.
- 8. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
- 9. Click OK.

Customizing the Aggregate Capacity and Utilization report to display aggregates with overcommitted threshold breached

You can customize the Aggregate Capacity and Utilization report to display the aggregates sorted by overcommitted capacity percentage, which enables you to view the storage space still available in the aggregates.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by cluster or HA pair, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To display the difference between the overcommitted used percentage and the overcommitted threshold, add a new column.
 - a. Select a column and click
 - b. Select Column > New Computed Column.
 - c. In the New Computed Column dialog box, enter a column label.
 - d. From the Select Category list, select **Math**.
 - e. From the **Select Function** list, select **DIFFERENCE**.
 - f. From the Column 1 list, select Overcommitted Threshold (%).
 - g. From the Column 2 list, select Overcommitted Capacity %.
 - h. Click OK.
- 3. To filter the values greater than zero in the new column, click in the **New computed column** and open the **Filter** dialog box by clicking the icon.
- 4. From the Condition list, select Greater Than.
- 5. In the Value field, type 0 and click OK.
- 6.
 To sort the values, click inside **New computed column** and click the icon.
- Select Filter > Top/Bottom N.
- 8. In the Top/Bottom N dialog box, select Top N from the Filter field and enter a value in the text field.
- 9. Click OK.

Customizing the Aggregate Capacity and Utilization report to display aggregates with noncompliant configuration

You can customize the Aggregate Capacity and Utilization report to display the aggregates filtered by the full threshold. This enables you to view the aggregates that might not comply with company policies.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

- 1. To remove the grouping by cluster or HA pair, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b.

Click the icon.

- c. Select Group > Delete Inner Group.
- 2. To filter aggregates threshold not exceeding 85%, click in the **Space Full Threshold** column and open the **Filter** dialog box by clicking the icon.
- 3. From the Condition list, select Greater Than.
- 4. Click Select Values and select 85.
- Click OK.

Volume Capacity and Utilization report customizations

You can create reports to monitor a variety of capacity and utilization information about volumes. For example, you can create reports to display volumes used, total capacity, daily growth rate, and Snapshot copy capacity, which can help you to determine if a volume is running out of space or whether it is being overutilized or underutilized.

Customizing the Volume Capacity and Utilization report to display volumes nearing full capacity with Snapshot Autodelete turned off

You can customize the Volume Capacity and Utilization report to display volumes sorted by increasing order of their volume capacity utilization. This enables you to display volumes reaching their full capacity.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

- 1. To remove the grouping by SVM, cluster, or volume, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2.

 To sort volumes that are nearing full capacity, click in the **Days To Full** column, and click the icon
- 3. To filter volumes that have Snapshot Autodelete turned off, click in the **Snapshot Autodelete** column and open the **Filter** dialog box by clicking the icon.
- 4. From the Condition list, select Equal To.
- 5. Click Select Values and select Disabled.
- 6. Click OK.

Customizing the Volume Capacity and Utilization report to display the least consumed volumes with thin provisioning disabled

You can customize the Volume Capacity and Utilization report to display volumes based on their volume consumption.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by SVM, cluster, or volume, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select **Group > Delete Inner Group**.
- To sort volumes based on percentage consumed, click in the Used Data % column, and click the icon
- To filter volumes with thin provisioning disabled, click in the Thin Provisioned column and open the Filter dialog box by clicking the icon.
- 4. From the Condition list, select Equal To.
- 5. Click Select Values and select No.
- 6. Click OK.

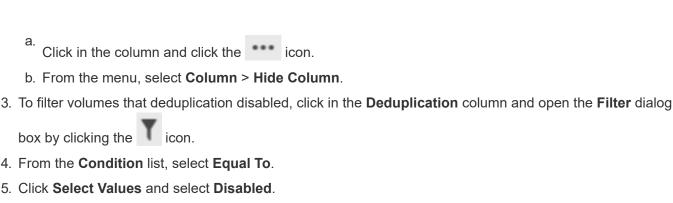
Customizing the Volume Capacity and Utilization report to display volumes with noncompliant configuration

You can customize the Volume Capacity and Utilization report to display volumes that are not compliant with company policies. For example, if you must have deduplication enabled on all volumes, you can create a report listing all volumes where deduplication is disabled.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

- 1. To remove the grouping by SVM, cluster, or volume, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. Hide all columns except for the Cluster, Storage Virtual Machine, Volume, Deduplication, and Deduplication Space Savings (GB) columns:



- 6. Click **OK**.
- 7. To sort volumes based on deduplication space savings, click in the **Deduplication Space Savings (GB)**column and click the

Qtree Capacity and Utilization report customizations

You can create customized reports to analyze capacity and utilization of the system's qtrees. For example, you can create reports to sort qtrees to determine whether any have breached the disk or file soft limit.

Customizing the Qtree Capacity and Utilization report to display qtrees that have breached the disk soft limit

You can customize the Qtree Capacity and Utilization report to display qtrees that have breached the disk soft limit. You can filter and sort by disk used, disk hard limit, and disk soft limit.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by SVM, cluster or volume, perform the following steps:
 - a. Click in the columns that need to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To filter gtrees that do not have an unlimited disk hard limit, click in the Disk Hard Limit column and open

the **Filter** dialog box by clicking the icon.

- a. From the Condition drop-down list, select Not Equal To.
- b. Click Select Values and select Unlimited.
- c. Click Ok.
- 3. To filter qtrees that do not have an unlimited disk soft limit, click in the **Disk Soft Limit** column and open the **Filter** dialog box by clicking the icon.

- a. From the Condition drop-down list, select Not Equal To.
- b. Click Select Values and select Unlimited.
- c. Click Ok.
- 4. To add a column for qtrees that have breached the disk soft limit, perform the following steps:
 - a.
 Click in the **Disk Soft Limit** column, click the icon, and select **Column > New Computed**Column
 - b. In the New Computed Column dialog box, type Breached Disk Soft Limit Capacity in the Column Label field.
 - c. From the Select Category list, select Text.
 - d. From the **Select Function** drop-down list, select **Advanced**.
 - e. In the Enter Expression field, type IF(([qtreeDiskUsedPercent] *[diskLimit]/100 >
 [softDiskLimit]), "Yes", "No").
 - f. Click OK.
- 5. To filter qtrees that have breached the soft disk limit, click in the **Breached Disk Soft Limit Capacity** column and open the **Filter** dialog box by clicking the icon.
 - a. From the Condition drop-down list, select Equal To.
 - b. Click Select Values and select Yes.
 - c. Click Ok.

Customizing the Qtree Capacity and Utilization report to display qtrees that have breached the file soft limit

You can customize the Qtree Capacity and Utilization report to display qtrees that have breached the file soft limit. You can filter and sort by file used, file hard limit, and file soft limit.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by SVM, cluster or volume, perform the following steps:
 - a. Click in the columns that need to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To filter qtrees that do not have an unlimited file hard limit, click in the File Hard Limit column and open the

Filter dialog box by clicking the icon.

- a. From the Condition drop-down list, select Not Equal To.
- b. Click Select Values and select Unlimited.

- c. Click Ok.
- 3. To filter qtrees that do not have an unlimited file soft limit, click in the File Soft Limit column and open the

Filter dialog box by clicking the icon.

- a. From the Condition drop-down list, select Not Equal To.
- b. Click **Select Values** and select **Unlimited**.
- c. Click Ok.
- 4. To add a column for qtrees that have breached the file soft limit, perform the following steps:
 - a.
 Click in the File Soft Limit column, click the column icon, and select Column > New Computed Column
 - b. In the New Computed Column dialog box, type Breached File Soft Limit Capacity in the Column Label field.
 - c. From the Select Category list, select **Text**.
 - d. From the **Select Function** drop-down list, select **Advanced**.
 - e. In the Enter Expression field, type IF(([qtreeFileUsedPercent]*[fileLimit]/100 > [softFileLimit]), "Yes", "No").
 - f. Click OK.
- 5. To filter qtrees that have breached the soft file limit, click in the **Breached File Soft Limit Capacity** column and open the **Filter** dialog box by clicking the icon.
 - a. From the Condition drop-down list, select Equal To.
 - b. Click Select Values and select Yes.
 - c. Click Ok.

Events report customizations

You can create reports to monitor outstanding events on a cluster.

Customizing the Events report to display events with a critical severity type

You can customize the Events report to display events filtered by their severity type, and by the events that have been unresolved for the longest period of time.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

- 1. To filter events with critical severity type, click in the **Status** column and open the **Filter** dialog box by clicking the icon.
- 2. From the Condition list, select Equal To.

- Click Select Values and select Critical.
- 4. Click OK.
- 5. To sort the events that are unresolved for the longest period of time, click in the **Days Outstanding** column, and click the icon.
- 6. Select Filter > Top/Bottom N.
- 7. In the Top/Bottom N dialog box, select Top N from the Filter field and enter a value in the text field.
- 8. Click OK.

Customizing the Events report to display events on mission-critical objects

You can customize the Events report to display events filtered by mission-critical data priority.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To filter events with mission-critical data priority, click in the **Data Priority** column and open the **Filter** dialog box by clicking the icon.
- 2. From the Condition list, select Equal To.
- 3. Click Select Values and select Mission-Critical.
- 4. Click OK.

Customizing the Events report to display the top most discussed events

You can customize the Events report to display events that are most discussed.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- To sort the events that are discussed the most, click in the **Notes** column and click the icon.
- Select Filter > Top/Bottom N.
- 3. In the **Top/Bottom N** dialog box, select **Top N** from the **Filter** field and enter a value in the text field.
- 4. Click OK.

Customizing the Events report to display incident events assigned to the admin

You can customize the Events report to display incident events that are assigned to the admin, filtered by the impact level and the admin name.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To filter incident events, click in the Impact Level column and open the Filter dialog box by clicking the
 - Yicon
- 2. From the Condition list, select Equal To.
- 3. Click Select Values and select Incident.
- Click OK.
- 5. To assign these incidents to the admin, click in the **Assigned To** column and open the **Filter** dialog box by clicking the icon.
- 6. From the Condition drop-down list, select Equal To.
- 7. Click Select Values and select Admin Name.
- 8. Click OK.

Customizing the Events report to display events impacting availability

You can customize the Events report to display events that are categorized by the most incidents and are assigned to the admin. You can filter the report by the impact level and the admin name.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To filter availability events, click in the **Impact Area** column and open the **Filter** dialog box by clicking the
 - Yicon
- 2. From the Condition drop-down list, select Equal To.
- 3. Click **Select Values** and select **Incident**.
- 4. Click OK.

Customizing the Events report to display the top most acknowledged unresolved events

You can customize the Events report to display the most acknowledged events, filtered by the event state. You can sort them in decreasing order to display the number of outstanding days.

About this task

You can also perform this task by going to the Reports page and clicking Run Report for the appropriate report.

Steps

1.

To filter acknowledged events, click in the **State** column and open the **Filter** dialog box by clicking the icon.



- 2. From the **Condition** drop-down list, select **Equal To**.
- Click Select Values and select Acknowledged.
- Click OK.
- 5. To further filter the report, click in the Acknowledged By column and open the Filter dialog box by clicking icon.
- 6. From the **Condition** drop-down list, select **Equal To**.
- 7. Click Select Values and select Name.
- 8. Click OK.
- 9. To sort the events that are outstanding for the most number of days, click in the **Days Outstanding** column and click
- 10. Select Filter > Top/Bottom N.
- 11. In the Top/Bottom N dialog box, select Top N from the Filter field and enter a value in the text field.
- 12. Click OK.

Cluster Inventory Report customizations

You can customize inventory reports to monitor for insufficient resources on clusters components. For example, you can customize reports to monitor information such as clusters that are nearing the SVM count limit, nodes that are running older versions of ONTAP, and nodes that are reaching the maximum disk limit.

Customizing the Cluster Inventory report to display clusters reaching SVM count limit

You can customize the Cluster Inventory report to display clusters, sorted by decreasing order of their SVM count.

About this task

You can also perform this task by going to the Reports page and clicking Run Report for the appropriate report.

Steps

- 1. To remove the grouping by cluster or node, perform the following steps:
 - a. Click in the column that needs to be ungrouped.

- b. Click the icon.
- c. Select Group > Delete Inner Group.
- 2. To sort clusters by SVM count, perform the following steps:
 - a. Click in the SVM Count column.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group option.
- 3. Select Filter > Top/Bottom N.
- 4. In the Top/Bottom N dialog box, select Top N from the Filter field and enter a value in the text field.
- 5. Click OK.

Customizing the Cluster Inventory report to display nodes running older versions of ONTAP software

You can customize the Cluster Inventory report to display nodes filtered by older ONTAP versions.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by cluster, or node, perform the following steps:
 - a. Click in the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To filter nodes not running ONTAP 8.3, click the **ONTAP version** column and open the **Filter** dialog box by clicking the icon.
- 3. From the Condition drop-down list, select Not Equal To.
- Click Select Values and select 8.3.
- 5. Click OK.

Customizing the Cluster Inventory report to display nodes reaching the maximum disk limit

You can customize the Cluster Inventory report to display a list of nodes that are reaching the maximum disk limit and sorted by increasing order.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by cluster, or node, perform the following steps:
 - a. Click in the columns that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To move the **Disk Count** column next to the **Model** column, perform the following steps:
 - a. Click in the Disk Count column.
 - b. Click the icon and select **Column > Reorder Columns**.
 - c. In the **Reorder Columns** dialog box, use the **up** and **down** arrow keys to move the column to the required position.
- 3. To add a new computed column, perform the following steps:
 - a. Select a column, click , and select Column > New Computed Column.
 - b. In the New Computed Column dialog box, type Maximum Disk Limit in the Column Label field.
 - c. From the Select Category list, select Comparison.
 - d. From the Select Function list, select Advanced.
 - e. In the Enter Expression field, type IF ([model]="FAS3250", 960, 0).
 - f. Click OK.
- 4. To add a second new column, perform the following steps:
 - a. Select the **Maximum Disk Limit** column, click the icon, and select **Column > New Computed Column**.
 - b. In the New Computed Column dialog box, type Available Volume in the Column Label field.
 - c. From the Select Category list, select Math.
 - d. From the **Select Function** list, select **DIFFERENCE**.
 - e. From the Column 1 list, select Maximum Disk Limit.
 - f. From the Column 2 list, select Disk Count.
 - g. Click OK.
- 5. To sort the values, click in the **Available Volume** column, and click the icon.
- Select Filter > Top/Bottom N.
- 7. In the Top/Bottom N dialog box, select Top N from the Filter field and enter a value in the text field.
- 8. Click OK.

NFS Export report customizations

You can customize NFS export reports to analyze information about NFS export policies and rules for volumes on your storage systems. For example, you can customize reports to display volumes with inaccessible junction paths and volumes with the default export policy.

Customizing the NFS Exports report to display a list of volumes that have an inaccessible junction path

You can customize the NFS Exports report to display a list of volumes that have an inaccessible junction path.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by cluster or volume, perform the following steps:
 - a. Click in the columns that need to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To filter volumes that have an inaccessible junction path, click in the **Junction Path Active** column and open the **Filter** dialog box by clicking the icon.
- 3. From the Condition list, select Equal To.
- 4. Click Select Values and select No.
- 5. Click OK.

Customizing the NFS Exports report to display a list of volumes with default export policy

You can customize the NFS Exports report to display a list of volumes with default export policy.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by cluster or volume, perform the following steps:
 - a. Click in the columns that need to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To filter volumes with default export policy, click the **Export Policy** column and open the **Filter** dialog box by clicking the icon.
- 3. From the **Condition** list, select **Equal To**.
- 4. Click Select Values and select Default.
- 5. Click OK.

SVM Inventory report customizations

You can create SVM inventory reports to analyze volume information and to view overall health and storage availability. For example, you can create reports to display SVMs reaching the maximum volume count and to analyze stopped SVMs.

Customizing the SVM Inventory report to display a list of SVMs reaching maximum volume limit

You can customize the SVM Inventory report to display a list of SVMs that are reaching the maximum volume limit by sorting the volumes in increasing order.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To remove the grouping by cluster, perform the following steps:
 - a. Click inside the column that needs to be ungrouped.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To filter SVMs that do not have unlimited allowed volumes, click the **Maximum Allowed Volumes** column and open the **Filter** dialog box by clicking the icon.
- 3. In the Data type field, select String and click OK.
- 4. From the Condition drop-down list, select Not Equal To.
- 5. Click **Select Values** and select **Unlimited**.
- 6. To add a new computed column, perform the following steps:
 - a. Select a column, click the icon, and select Column > New Computed Column.
 - b. In the New Computed Column dialog box, type Available Volume in the Column Label field.
 - c. From the Select Category list, select **Math**.
 - d. From the **Select Function** drop-down list, select **Advanced**.
 - e. In the **Enter Expression** field, type [maximumVolumes] [volumeCount].
 - f. Click OK.
- To sort SVMs in ascending order, click in the **Available Volume** column, and click the icon.
- 8. Select Filter > Top/Bottom N.
- 9. In the **Top/Bottom N** dialog box, select **Bottom N** from the **Filter** field and enter a value in the text field.
- 10. Click **OK**.

Customizing the SVM Inventory report to display a list of stopped SVMs

You can customize the SVM Inventory report to display a list of stopped SVMs. The report filters the SVMs by their status.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- To filter SVMs by status, click the **State** column and open the **Filter** dialog box by clicking the icon.
- 2. From the Condition list, select Equal To.
- Click Select Values and select Stopped.
- 4. Click OK.

Volume Relationships Inventory report customizations

You can customize the Volume Relationships Inventory report to view the volume details that are filtered based on the source of failure. You can use filters to display volume relationships inventory details based on schedules, and to group volume inventory details based on issues.

Customizing the Volume Relationships Inventory report to view volumes grouped by source of failure

You can customize the Volume Relationships Inventory report to view volumes grouped by source of failure.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. Select the **Relationship Health** column.
- 2. To view the volume details for bad volumes, click the sign next to the **Bad** column.
- 3. To view the volume details for good volumes, click the sign next to the **Good** column.

Customizing the Volume Relationships Inventory report to view volumes grouped by issue

You can customize the Volume Relationships Inventory report to view volumes that are grouped according to the volume relationship health status.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. To filter volumes according to the volume relationship health status, select the **Relationship Health** column, and click the icon.
- 2. In the **Filter** dialog box, click **Select Values**, and then select the required value from the drop-down list.

The volume details for the selected value are displayed.

Volume Transfer Status (Historical) report customizations

You can customize the Volume Transfer Status (Historical) report to view and analyze information about volume transfers at specific time intervals. You can use filters to view volume transfer details between two dates.

Customizing the Volume Transfer Status (Historical) report schedules

You can customize the schedules for the Volume Transfer Status (Historical) report to view the volume details based on different schedules. You can view, modify, or delete existing report schedules, and add new schedules for your reports.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. On the Volume Transfer Status (Historical) report page, click Manage Report Schedules.
- 2. In the **Manage Report Schedules** dialog box, enter specific details such as recipient schedule name, email address, report format, frequency, and the reports.
- Select Inventory as the Report Category.
- Click Save and Close.

The Volume Transfer Status (Historical) report is automatically sent by email to one or more recipients as per the schedule.

Customizing the Volume Transfer Status (Historical) report to view volumes at specific time intervals

You can customize the Volume Transfer Status (Historical) report to view the volume details at specific time intervals.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

1. Remove grouping by cluster:

- a. Click in the column that you want to ungroup.
- b. Click the icon.
- c. Select Group > Delete Inner Group.
- 2. To view the volume details at a specific time interval, click in the **Start time** column, and then click the icon.
- In the Filter dialog box, click Select Values, and then select the specific date and time from the drop-down list.

The volume details for the selected time range are displayed.

Customizing the Volume Transfer Status (Historical) report to view volumes grouped by time of occurrence

You can customize the Volume Transfer Status (Historical) report to display the list of volumes grouped by time of occurrence between two dates.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. Remove grouping by cluster:
 - a. In the column that has to be ungrouped, click the icon.
 - b. Select **Group > Delete Inner Group**.
- 2. In the **Start time** column, open the **Filter** dialog box by clicking the icon.
- 3. From the Condition drop-down list, select Between.
- 4. Click **Select Values**, and choose the **Date From** and **Date To** values.
- 5. Click OK.

Customizing the Volume Transfer Status (Historical) report to view failed or successful volume transfers

You can customize the Volume Transfer Status (Historical) report to view the details of failed or successful volume transfers.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. Remove grouping by cluster:
 - a. Select the column that you want to ungroup.

- b. Click the icon
- c. Select Group > Delete Inner Group.
- 2. To sort the volume transfers according to failure or success, click in the **Operational Result** column, and then click the icon.
- 3. Select Filter.
- 4. In the Filter dialog box, click Select Values, and then select either Success or Failure.

Volume Transfer Rate (Historical) report customizations

You can customize the Volume Transfer Rate (Historical) report to view the volume transfer details based on the total transfer size of the volume. You can also view the volume transfers for a specific day of the week.

Customizing the Volume Transfer Rate (Historical) report to view volume transfers based on transfer size

You can customize the Volume Transfer Rate (Historical) report to view the volume transfer details according to the total transfer size of the volume.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. Remove grouping by cluster:
 - a. Select the column that you want to ungroup.
 - b. Click the icon.
 - c. Select Group > Delete Inner Group.
- 2. To sort the volume transfers according to volume transfer size, click the **Total Transfer Size (GB)** column.

Customizing the Volume Transfer Rate (Historical) report to view volume transfers grouped by day

You can customize the Volume Transfer Rate (Historical) report to view the volume transfer details that are sorted by day.

About this task

You can also perform this task by going to the Reports page and clicking **Run Report** for the appropriate report.

Steps

- 1. Remove grouping by cluster:
 - a. Select the column that you want to ungroup.
 - b. Click the icon.

- c. Select **Group > Delete Inner Group**.
- 2. To view the volume transfers for a specific day, click the **Day** column.

Description of report windows and dialog boxes

You can use the options to schedule, share, manage, save, and import the reports.

Reports page

The Reports page enables you to view detailed information about the reports that you generate. You can search for a specific report, save a report, and delete a report. You can also schedule, share, and import a report.

The Reports page displays categorized groups of reports about which you can obtain specific report details. By default, the report groups expand to display the report types, a report overview, and links that enable you to customize reports. Only one report can be viewed at a time. You can click the **Run Report** button to view a report for a specific group of reports.

The implementation of scheduling and generating reports from the Reports page will be deprecated in a future release. You can extract health and performance data from Unified Manager using these additional methods:



- Extracting data directly from the grid pages in the GUI
- Using Open Database for access to all the available objects
- Executing Unified Manager REST APIs to return the information you are interested in reviewing

The following is a list of report groups and report types that are displayed in the Reports page:

- · Capacity Utilization Reports
 - Storage Summary
 - Aggregate Capacity and Utilization
 - Volume Capacity and Utilization
 - Qtree Capacity and Utilization
- · Operational Reports Events
- · Inventory Reports
 - Cluster Inventory
 - NFS Exports
 - SVM Inventory
- · Imported Reports
- Data Protection Reports
 - Volume Data Protection Configuration
 - Volume Relationships Inventory
 - Volume Transfer Status (Historical)

Volume Transfer Rate (Historical)

Storage Summary report

The Storage Summary report enables you to view summarized information about storage capacity in the HA pairs. This information helps you to understand possible capacity risks and to take appropriate action to rebalance workload. Single-node cluster information is not visible in the report.

Storage Summary report chart view

The Storage Summary report is displayed in two formats:

- · Storage Summary report chart view
- Storage Summary report tabular view

The chart shows the capacity trend of used and unused data capacity of the aggregates over a period of time. Total data capacity is displayed on the vertical (y) axis and the cluster name on the horizontal (x) axis. Therefore, each bar in the chart represents one cluster. You can view the details for specific points on the graph by positioning your cursor over a particular point.

Cluster Name

Displays the cluster name.

HA Pair

Displays the HA pair value obtained by forming two nodes.

Model

Displays the name of the model.

OS Version

Displays the version of ONTAP used.

Total Raw Capacity

Displays the total physical capacity of all disks in the array.

Unconfigured Raw Capacity

Displays the unconfigured capacity of disks whose container type is other than aggregate, broken, spare, or shared. This capacity is always higher than the physical capacity of the disk in ONTAP. For example, consider a 2 TB disk. The physical capacity of the disk is 1.6 TB in ONTAP whereas the unconfigured raw capacity in Unified Manager is 1.8 TB.

Aggregate Total Capacity

Displays the total size of the available aggregates for the user. This includes the Snapshot copy reserve.

Aggregate Used Capacity

Displays the capacity already in use on aggregates. This includes the capacity consumed by volumes, LUNs, and other storage efficiency technology overheads.

Aggregate Unused Capacity

Displays capacity that might be available for storing additional data on the aggregate. This includes the Snapshot copy reserve.

Allocated LUN Capacity

Displays the capacity of LUNs that are mapped.

Unallocated LUN Capacity

Displays the capacity of all LUNs not mapped to the Host.

Volume Total Capacity

Displays the total capacity of the volumes (used plus unused).

Volume Used Capacity

Displays the used capacity of the volumes.

Volume Unused Capacity

Displays the unused capacity of the volumes.

Volume Protection Capacity

Displays the capacity of volumes that have SnapMirror and SnapVault enabled.

Cluster Licensed Cloud Tier Total

Displays the total capacity that has been licensed in the cloud tier. This field is displayed for storage providers that require a FabricPool license, for example, Amazon S3, IBM Cloud Object Storage, Microsoft Azure Cloud, or Alibaba Cloud Object Storage.

Cluster Licensed Cloud Tier Used

Displays the space used by data in the cloud tier for storage providers that require a FabricPool license.

Cluster StorageGRID Capacity Used

Displays the space used by data in the cloud tier for storage providers that do not require a FabricPool license, for example, StorageGRID.

Aggregate Capacity and Utilization report

The Aggregate Capacity and Utilization report enables you to view information about the capacity and utilization of aggregates in a cluster. This information enables you to understand possible capacity risks and also to view the configured, used, and unused capacity of aggregates.

Aggregate Capacity and Utilization report tabular view

Cluster

Displays the cluster name.

• HA Pair

Displays the HA pair value obtained by forming two nodes.

Aggregate

Displays the aggregate name.

Total Data Capacity (GB)

Displays the total data capacity (used plus available).

Used Data Capacity (GB)

Displays the used data capacity.

Used Data %

Displays the used data capacity as a percentage.

Available Data Capacity (GB)

Displays the available data capacity.

Available Data %

Displays the available data capacity as a percentage.

Daily Growth Rate %

Displays the growth rate that occurs every 24 hours in the volume.

Days To Full

Displays the estimated number of days remaining before the aggregate reaches full capacity.

Space Full Threshold

Displays the percentage at which an aggregate is full.

Space Nearly Full Threshold

Displays the percentage at which an aggregate is nearly full.

· Growth Rate Threshold

Specifies the aggregate's growth rate is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event.

Growth Rate Sensitivity Threshold

Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated.

Days Until Full Threshold

Specifies the number of days remaining before the aggregate reaches full capacity.

Snapshot Reserve Total Capacity (GB)

Displays the total snapshot reserve capacity of the aggregate.

Snapshot Reserve Used Capacity (GB)

Displays the amount of space used by snapshot copies from snapshot reserve.

Snapshot Reserve Used %

Displays the amount of space used by Snapshot copies from snapshot reserve as a percentage.

Snapshot Reserve Available Capacity (GB)

Displays the amount of space available for Snapshot copies.

Snapshot Reserve Available %

Displays the amount of space available for Snapshot copies as a percentage.

· Snapshot Copies Reserve Full Threshold

Specifies the percentage at which an aggregate has consumed all its space reserved for Snapshot copies.

Overcommitted Capacity %

Displays the aggregate overcommitment as a percentage.

Overcommitted Threshold %

Displays the percentage at which an aggregate is overcommitted.

Nearly Overcommitted Threshold %

Displays the percentage at which an aggregate is nearly overcommitted.

Type

Displays the aggregate type:

- HDD
- Hybrid

Combines HDDs and SSDs, but Flash Pool has not been enabled.

Hybrid (Flash Pool)

Combines HDDs and SSDs, and Flash Pool has been enabled.

- · SSD
- SSD (FabricPool)

Combines SSDs and a cloud tier

VMDisk (SDS)

Virtual disks within a virtual machine

VMDisk (FabricPool)

Combines virtual disks and a cloud tier

LUN (FlexArray)

For standard disks and SSD disks, this column is blank when the monitored storage system is running an ONTAP version earlier than 8.3.

RAID Type

Displays the RAID configuration type.

Aggregate State

Displays the current state of the aggregate.

SnapLock Type

Indicates whether the aggregate is a SnapLock or non-SnapLock aggregate.

Cloud Tier Space Used (GB)

Displays the amount of data capacity that is currently being used in the cloud tier.

Cloud Tier

Displays the name of the cloud tier when it was created by ONTAP.

Volume Capacity and Utilization report

The Volume Capacity and Utilization report enables you to view information about the capacity and utilization of volumes in a cluster. This information enables you to understand possible capacity risks and to view the configured, used, and unused capacity of aggregates. Also, the report helps you to make decisions about enabling space-saving features such as deduplication and thin provisioning.

Volume Capacity and Utilization report tabular view

Cluster

Displays the cluster name.

Storage Virtual Machine

Displays the name of the storage virtual machine (SVM) that contains the volume.

Volume

Displays the volume name.

Total Data Capacity

Displays the total data capacity (used plus available) in a volume.

Used Data Capacity

Displays the used data capacity in a volume.

Used Data %

Displays the used data in a volume as a percentage.

Available Data Capacity

Displays the available data capacity in a volume.

Available Data %

Displays the available data capacity in a volume as a percentage.

Daily Growth Rate %

Displays the growth rate that occurs every 24 hours in the volume.

Days To Full

Displays the estimated number of days remaining before the volume reaches full capacity.

Space Full Threshold %

Specifies the limit to the volume that is considered full.

Space Nearly Full Threshold %

Specifies the limit to the volume that is considered nearly full.

Growth Rate Threshold %

Specifies the aggregate's growth rate is considered to be normal before the system generates an Aggregate Growth Rate Abnormal event.

Growth Rate Sensitivity Threshold

Specifies the factor that is applied to the standard deviation of a volume's growth rate. If the growth rate exceeds the factored standard deviation, a Volume Growth Rate Abnormal event is generated.

Days Until Full Threshold

Specifies the number of days remaining before reaching full capacity.

Snapshot Overflow %

Displays the percentage of the data space that is consumed by the Snapshot copies.

Snapshot Reserve Used Capacity

Displays the amount of space used by Snapshot copies in the volume.

Snapshot Reserve Used %

Displays the amount of space used by Snapshot copies in the volume as a percentage.

Snapshot Reserve Available Capacity

Displays the amount of space available for Snapshot copies in the volume.

Snapshot Reserve Available %

Displays the amount of space available for Snapshot copies in the volume as a percentage.

Snapshot Reserve Total Capacity

Displays the total Snapshot copy capacity in the volume.

Snapshot Copies Reserve Full Threshold %

Specifies the percentage at which the space reserved for Snapshot copies is considered full.

Snapshot Copies Count Threshold

Specifies the number of Snapshot copies on a volume that are considered to be too many.

Snapshot Copies Days Until Full Threshold

Specifies the number of days remaining before the space reserved for Snapshot copies reaches full capacity.

Number Of Inodes

Displays the number of inodes in the volume.

Inode Utilization %

Specifies the percentage of inode space used in the volume.

Inodes Full Threshold

Specifies the percentage at which a volume is considered to have consumed all of its inodes.

Inodes Nearly Full Threshold

Specifies the percentage at which a volume is considered to have consumed most of its inodes.

Quota Committed Capacity

Displays the space reserved in the volumes.

Quota Overcommitted Capacity

Displays the amount of space that can be used before the system generates the Volume Quota Overcommitted event.

Quota Overcommitted Threshold %

Specifies the percentage at which the volume is nearly overcommitted.

Quota Nearly Overcommitted Threshold %

Specifies the percentage at which the volume space is nearly overcommitted.

Snapshot Autodelete

Displays whether automatic deletion of Snapshot copies is enabled or disabled.

Deduplication

Displays whether deduplication is enabled or disabled for the volume.

Deduplication Space Savings

Displays the amount of space saved in a volume by using deduplication.

Compression

Displays whether compression is enabled or disabled for the volume.

Compression Space Savings

Displays the amount of space saved in a volume by using compression.

Caching Policy

Displays the caching policy that is associated with the selected volume. The policy provides information about how Flash Pool caching occurs for the volume. See the Health/Volumes inventory page for more information on caching policies.

Cache Retention Priority

Displays the priority used for retaining cached pools.

Thin Provisioned

Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.

Autogrow

Displays whether the FlexVol volume automatically grows in size when it is out of space.

Space Guarantee

Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate.

State

Displays the state of the volume that is being exported.

SnapLock Type

Indicates whether the volume is a SnapLock or non-SnapLock volume.

Expiry Date

The SnapLock expiration date.

Tiering Policy

If this volume is deployed on a FabricPool-enabled aggregate, then the tiering policy set for the volume is displayed.

Qtree Capacity and Utilization report

The Qtree Capacity and Utilization report enables you to analyze capacity and utilization of the system's qtrees to understand possible risks that might occur due to reduced cluster capacity.

Qtree Capacity and Utilization report tabular view

Cluster

Displays the name of the cluster containing the qtree.

Storage Virtual Machine

Displays the storage virtual machine (SVM) name containing the qtree.

Volume

Displays the name of the volume containing the qtree.

Qtree

Displays the name of the qtree.

Quota type

Specifies if the quota is for a user, user group or a qtree.

User or Group

Displays the name of the user or user group. There will be multiple rows for each user and user group. When the quota type is qtree, then *Not Applicable* is displayed. If the quota is not set, then the column is empty.

Disk Used %

Displays the percentage of the disk space used. If a disk hard limit is set, this value is based on the disk hard limit. If the quota is set without a disk hard limit, the value is based on the volume data space. If the quota is not set or if the quotas are off on the volume to which the qtree belongs, then *Not applicable* is displayed.

Disk Hard Limit

Displays the maximum disk space allocated for the qtree. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as *Unlimited* if the quota is set without a disk hard limit, If the quota is not set, or if the quotas are off on the volume to which the qtree belongs.

Disk Soft Limit

Displays the disk space allocated for the qtree before a warning event is generated. The value is displayed as *Unlimited* if the quota is set without a disk soft limit, if the quota is not set, or if the quotas are off on the volume to which the qtree belongs.

Files Used %

Displays the percentage of files used in the qtree. If the file hard limit is set, this value is based on the file hard limit. The value is displayed as *Not applicable* if the quota is not set, or if the quota is set without a file hard limit, or if the quotas are off on the volume to which qtree belongs.

File Hard Limit

Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as *Unlimited* if the quota is set without a file hard limit, if the quota is not set, or if the quotas are off on the volume to which the qtree belongs.

File Soft Limit

Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as *Unlimited* if the quota is set without a file soft limit, if the quota is not set, or if the quotas are off on the volume to which the qtree belongs.

Events report

The Events report enables you to view information about event trends over a specific time period. This information enables you to compare recent activity with any past operational activity, such as configuration changes, upgrades, and so on. The information also helps you to determine any outstanding events.

Events report chart view

The Events report is displayed in two formats:

- Events report chart view
- Events report tabular view

The Events Chart is displayed in two formats:

- Events Severity Trend (All open events)
- · Event Status Trend

The chart shows the event severity trends for all open events over a time period. A count of events is displayed on the vertical (y) axis and the date is displayed on the horizontal (x) axis. You can view the details for specific points on the graph by positioning your cursor over a particular point. The details display the event severity,

number of events of the specific severity type, and the date of the event.

The event severity types displayed are Critical, Error, and Warning. The event severities are differentiated by different colors. There can be the same number of events on the same date in different states.

Count

Displays a count of events.

Date

Displays the date. The x axis shows data from the time that the event occurred up to the present date. You can click and zoom the chart to get details.

The chart shows the event status trending per day over a period of time. A count of events is displayed on the vertical (y) axis and the date is displayed on the horizontal (x) axis. The details display the event state, number of events of the specific state, and the date of the event.

The event status are New, Acknowledged, and Resolved. The event status are differentiated by different colors.

The chart shows the new events generated daily on a cumulative basis in a bar graph represented in green color. The number of Acknowledged and Resolved events are shown as and when they are acknowledged and resolved on a daily basis.

There is a zoom functionality provided within the charts. You can use this feature to zoom a particular point in the chart for more clarity.

Source

Displays the source of an event.

Status

Displays the severity of the event. You can filter this column to display events of a specific severity type. The event severity types are Critical, Error, or Warning.

State

Displays the event state: New, Acknowledged, Resolved, or Obsolete. You can filter this column to show events of a specific state.

Event

Displays the event names.

Triggered Time

Displays the time when the event was generated. Both the time and the date are displayed.

Days Outstanding

Displays the number of days between an event occurring and its resolution or designation as Obsolete.

Source Type

Displays the object type (for example, Storage Virtual Machine (SVM), volume, or qtree) with which the event is associated.

Data Priority

Displays the annotation type, based on the priority of data of the storage object.

Impact Level

Displays whether the event is categorized as an incident, a risk, or information.

Impact Area

Displays whether the event is a capacity, availability, performance, protection, or configuration event.

Assigned To

Displays the name of the user to whom the event is assigned.

Assigned Time

Displays the time when the event was assigned to a user.

Notes

Displays the number of notes that are added for an event.

· Acknowledged By

Displays the name of the user who acknowledged the event. The field is blank if the event is not acknowledged.

Acknowledged Time

Displays the time that has elapsed since the event was acknowledged. If the time elapsed exceeds a week, the timestamp displays when the event was acknowledged.

Resolved By

Displays the name of the user who resolved the event. The field is blank if the event is not resolved.

Resolved Time

Displays the time that has elapsed since the event was resolved. If the time elapsed exceeds a week, the timestamp displays when the event was resolved.

Obsoleted Time

Displays the time when the state of the event became Obsolete.

Cluster Inventory report

Cluster Inventory report provides information about available resources for cluster components for the purpose of understanding possible risks caused by insufficient resources.

Cluster Inventory report tabular view

Cluster

Displays the name of the cluster.

• HA pair

Displays the HA pair value obtained by forming two nodes.

Node

Displays the name of the nodes.

Model

Displays the name of the model.

OS version

Displays the version of ONTAP used.

All Flash Optimized

Displays whether node is configured to support only solid-state drives (SSDs).

Serial Number

Displays the serial number of the node.

Firmware Version

Displays the firmware version of the node.

SVM Count

Displays the number of SVM contained by the cluster.

• FC Port Count

Displays the number of FC ports contained by the node.

FCoE Port Count

Displays the number of FCoE ports contained by the node.

Ethernet Port Count

Displays the number of ethernet ports contained by the node.

Flash Card Count

Displays the number of flash cards installed on nodes in your data center so that you can monitor for potential problems.

• Flash Card Size (GB)

Displays the size of the flash cards installed on nodes.

Disk Shelves Count

Displays the number of disk shelves contained by the node.

Disk Count

Displays the number of disks in a node.

NFS Exports report

NFS Exports report enables you to audit information about NFS export policies and its associated rules for volumes in your storage system.

NFS Exports report tabular view

Cluster

Displays the name of the cluster.

Storage Virtual Machine

Displays the name of the SVM with NFS export policies.

Volume

Displays the name of the volume with NFS export policies.

Qtree

Displays the name of the qtree on a volume with NFS export policies.

Volume State

Displays the current state of the volume. The state can be Offline, Online, or Restricted.

Offline

Read or write access to the volume is not allowed.

Online

Read and write access to the volume is allowed.

· Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Junction Path

Displays the path on which the volume is mounted.

Junction Path Active

Displays whether the path to access the mounted volume is active or inactive.

Export policy

Displays the rules that define the access permission for volumes that are exported.

Rule Index

Displays the rules associated with the export policy such as the authentication protocols and the access permission.

Access Protocols

Displays the protocols that are enabled for the export policy rules.

Client Match

Displays the clients that have permission to access data on the volumes.

Read Only Access

Displays the authentication protocol used to read data on the volumes.

Read Write Access

Displays the authentication protocol used to read or write data on the volumes.

Security Style

Displays the access permission for volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

UNIX (NFS clients)

Files and directories in the volume have UNIX permissions.

Unified

Files and directories in the volume have a unified security style.

NTFS (CIFS clients)

Files and directories in the volume have Windows NTFS permissions.

Mixed

Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

Unix Permission

Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.

SVM Inventory report

SVM Inventory report enables you to analyze SVM volume configuration limits and overall health to understand risks to future storage availability.

SVM Inventory report tabular view

Cluster

Displays the name of the cluster containing the SVM.

Storage Virtual Machine

Displays the name of the SVM.

State

Displays the current administrative state of the SVM. The state can be Running, Stopped, Starting, Stopping, Not mapped, Initializing, or Deleting.

Volume count

Displays the number of volumes contained by the SVM.

Maximum Allowed Volumes

Displays the maximum allowed volumes that can be configured on the SVM.

Root Volume

Displays the name of the root volume of the SVM.

Allowed protocols

Displays the type of protocols that can be configured on the SVM.

DNS Domain

Displays the DNS domain name.

NIS Domain

Displays the Network Information Service (NIS) domain name. This column is blank when the Network Information Service (NIS) server is disabled or is not configured.

LDAP Enabled

Displays if the LDAP protocol is enabled or not.

Name Service Switch

Displays the information type gathered from hosts. Possible values are file, LDAP, or NIS.

Volume Data Protection Configuration report

The Volume Data Protection Configuration report enables you to view the unprotected volumes and storage virtual machines (SVMs) that are used in a node or a cluster. This information enables you to understand the data protection risks for your system, and to view the details of the protected volumes and unprotected volumes in your system.

The Volume Data Protection Configuration report is displayed in two formats:

- · Protected and Unprotected Volumes pie chart
- · Unprotected Volume Data tabular view

Protected and Unprotected Volumes pie chart

Displays the relative percentage of the protected volumes and unprotected volumes in your system.

Unprotected Volume Data tabular view

Cluster

Displays the cluster name.

Storage Virtual Machine

Displays the name of the storage virtual machine (SVM) that contains the volume.

Volume

Displays the volume name.

Total Data Capacity (GB)

Displays the total data capacity (used plus available) in GB.

Used Data Capacity (GB)

Displays the used data capacity (in GB).

Used Data %

Displays the used data capacity as a percentage.

Available Data Capacity (GB)

Displays the available data capacity (in GB).

Available Data %

Displays the available data capacity as a percentage.

Snapshot Reserve Used Capacity (GB)

Displays the amount of space that is used by Snapshot copies from Snapshot reserve (in GB).

Snapshot Reserve Used %

Displays the amount of space that is used by Snapshot copies from Snapshot reserve as a percentage.

Snapshot Reserve Available Capacity (GB)

Displays the amount of space that is available for Snapshot copies (in GB).

Snapshot Reserve Available %

Displays the amount of space that is available for Snapshot copies as a percentage.

Snapshot Reserve Total Capacity (GB)

Displays the total snapshot reserve capacity of the aggregate (in GB).

Days To Full

Displays the estimated number of days remaining before the aggregate reaches full capacity.

Space Full Threshold %

Displays the percentage at which an aggregate is full.

Space Nearly Full Threshold %

Displays the percentage at which an aggregate is nearly full.

Daily Growth Rate %

Displays the growth rate that occurs every 24 hours in the volume.

Total Number Of Inodes

Displays the total number of inodes in the volume.

Inode Utilization

Specifies the inode space that is used in the volume.

Quota Committed Capacity

Displays the space that is reserved in the volumes.

Quota Overcommitted Capacity (GB)

Displays the amount of space that can be used (in GB) before the system generates the Volume Quota Overcommitted event.

Snapshot Autodelete

Displays whether automatic deletion of Snapshot copies is enabled or disabled.

Deduplication

Displays whether deduplication is enabled or disabled for the volume.

Deduplication Space Savings (GB)

Displays the amount of space that is saved in a volume by using deduplication (in GB).

Compression

Displays whether compression is enabled or disabled for the volume.

Compression Space Savings (GB)

Displays the amount of space that is saved in a volume by using compression (in GB).

Thin Provisioned

Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.

Autogrow

Displays whether the FlexVol volume automatically grows in size when it is out of space.

Space Guarantee

Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate.

State

Displays the state of the volume that is being exported.

SnapLock Type

Indicates whether the volume is a SnapLock or non-SnapLock volume.

Expiry Date

Volume Relationships Inventory report

The Volume Relationships Inventory report enables you to analyze the storage inventory details in a cluster, understand the degree of protection that is required for volumes, and filter the volume details based on source of failure, pattern, and schedules.

The Volume Relationships Inventory report is displayed in two formats:

- SnapMirror relationships pie chart and SnapVault relationships pie chart
- · Volume Relationships Inventory report tabular view

SnapMirror and SnapVault pie charts

Displays the configuration details of the volume relationships that are present in your storage system.

Volume Relationships Inventory tabular view

· Relationship Health

Displays the relationship heath of the cluster.

Relationship State

Displays the the mirror state of the SnapMirror relationship.

Transfer Status

Displays the status of the SnapMirror relationship.

Lag Status

Displays the lag status of the volume.

Source Cluster

Displays the name of the source cluster for the SnapMirror relationship.

Source SVM

Displays the name of the source storage virtual machine (SVM) for the SnapMirror relationship.

Source Volume

Displays the name of the source volume for the SnapMirror relationship.

Destination Cluster

Displays the name of the destination cluster for the SnapMirror relationship.

Destination SVM

Displays the name of the destination storage virtual machine (SVM) for the SnapMirror relationship.

Destination Volume

Displays the name of the destination volume for the SnapMirror relationship.

Relationship Type

Displays any relationship type, including SnapMirror or SnapVault.

Last Successful Update Time

Displays the time of the last successful SnapMirror or SnapVault operation.

Last Transfer Duration (hrs)

Displays the time taken for the last data transfer to complete.

Last Transfer Size (MB)

Displays the size, in bytes, of the last data transfer.

Last Transfer End Time

Displays the time that the last successful SnapMirror or SnapVault operation completed.

Unhealthy Reason

The reason the relationship is in an unhealthy state.

Lag Duration (hrs)

Displays the amount of time that the data on the mirror lags behind the source.

Version Flexible Replication

Displays either Yes, Yes with backup option, or None.

Volume Transfer Status (Historical) report

The Volume Transfer Status (Historical) report enables you to analyze the volume transfer trends over a period of time. You can configure the report to view the volume transfer status for a specific time interval. The report also displays whether the volume transfer was a success or a failure.

The Volume Transfer Status (Historical) report is displayed in two formats:

- Volume Transfer Status line chart
- · Volume Transfer Status (Historical) report tabular view

Volume Transfer Status line chart

The line chart displays the volume transfer details by plotting transfer count against date. You can also view whether a particular volume transfer has succeeded or failed.

Volume Transfer Status tabular view

Source Cluster Name

Displays the source cluster name.

Source SVM

Displays the storage virtual machine (SVM) name.

Source Volume Name

Displays the source volume name.

Destination Cluster Name

Displays the destination cluster name.

Destination SVM

Displays the destination SVM name.

Destination Volume Name

Displays the destination volume name.

Operation Result

Displays whether volume transfer was successful.

Start time

Displays the volume transfer start time.

End time

Displays the volume transfer end time.

Transfer duration (hh:mm:ss)

Displays the time taken (in hours) to complete the volume transfer.

Transfer size (MB)

Displays the size (in MB) of the transferred volume.

Operation Type

Displays the type of volume transfer.

Volume Transfer Rate (Historical) report

The Volume Transfer Rate (Historical) report enables you to analyze the amount of data volume that is transferred on a day-to-day basis. The report also provides details about daily volume transfers and the time required to complete the transfer operation.

The Volume Transfer Rate (Historical) report is displayed in two formats:

- Volume Transfer Rate bar chart
- · Volume Transfer Rate tabular view

Volume Transfer Rate bar chart

Displays the volume transfer rate details by plotting the total transfer size against the number of hours. You can also view the details of the amount of data that is transferred on a daily basis.

Volume Transfer Rate tabular view

Total Transfer Size (GB)

Displays the total size of the volume transfer in gigabytes.

Day

Displays the day on which the volume transfer was initiated.

End Time

Displays the volume transfer end time with date.

Schedule Report dialog box

You can schedule the reports to be generated on a recurring basis at a specified frequency from the Schedule Report dialog box. The report is sent by email to one or more users specified in the Schedule Report dialog box.

Properties

You can schedule a report by specifying properties such as the email address of the user, the format of the report, and the frequency at which the report is generated.

Using Existing Schedule

Schedule Name

Displays all the existing schedule names. You can select an existing schedule for your reports from here.

Create New Schedule

Schedule Name

Enables you to enter the schedule name while creating a new schedule.

Recipient Email Address

Specifies the email address of the user to whom you want to send the report. You can specify one or more entries, separated by commas. This is a mandatory field.

Report Format

Specifies the format in which you want to schedule the report. The PDF option is selected by default.

Frequency

Specifies the frequency at which you want to schedule the report. The *Hourly* option is selected by default.

Command buttons

The command buttons enable you to perform the following tasks:

Schedule

Schedules the report with the saved or updated template and closes the Schedule Report dialog box.

Cancel

Closes the Schedule Report dialog box while displaying a message to save the schedule report template.

Share Report dialog box

You can share a report with one or more users through email. After you customize a report, you must save the changes before you share the report to ensure that the changes are displayed.

Properties

You can share a report by specifying properties such as the email address of the user, subject of the email, and the format of the report.

Recipient Email Address

Specifies the email address of the user with whom you want to share the report. You can specify one or more entries, separated by commas. This is a mandatory field.

Subject

Specifies the subject of the email. By default, the name of the report is displayed.

Report Format

Specifies the format in which you want to share the report. The *PDF* option is selected by default. If the XHTML format is selected, open the report that is sent by email by using a supported web browser.

Command buttons

The command buttons enable you to perform the following tasks:

Share

Shares the report with the saved configuration and closes the Share Report dialog box.

Cancel

Closes the Share Report dialog box while displaying a message to save the report configuration.

Manage Report Schedules dialog box

You can view, modify, or delete existing report schedules and add new schedules for your reports from the Manage Report Schedules dialog box.

Properties

You can select an existing schedule or create a new schedule for your reports. You can view, modify, or delete your report schedules.

· Left pane

Schedule Name

Displays the existing schedules. By clicking on any schedule you can view the schedule details in the right pane. For the first login, there are no existing schedules.

Add Schedule

Displays the new schedule form in the right pane. You can now add a new schedule.

· Right pane

Schedule Name

Displays the schedule name.

Recipient Email Address

Displays the email address of the user to whom the report must be sent. You can enter more than one email addresses separated by commas.

Report Format

Displays the format in which the report must be presented. The PDF option is selected as the default report format. If the XHTML format is selected, open the report that is sent by email by using a supported web browser.

Frequency

Displays the frequency at which the report is scheduled.

Report Category

Displays the report category groups. Selecting a report category from the list, displays the reports that belong to that report category in the Available Reports column.

Available Reports

Displays only the reports that belong to the report category selected.

Selected Reports

Displays the selected reports to which you choose to apply the schedule. You can select the required reports from the Available Reports column. At least one report must be selected

Command buttons

The command buttons enable you to perform the following tasks:

Add Schedule

Enables you to add a new schedule.

· Delete Schedule

Enables you to delete the schedule being currently viewed. When you create a new schedule, this button is not available.

Save

Saves the schedule being viewed, modified, or added.

Save and Close

Saves the schedule being viewed, modified, or added and closes the Manage Report Schedules dialog box.

Cancel

Closes the Manage Report Schedules dialog box while displaying a message to save the schedule.

Save Customized Report As dialog box

You can use the Save Customized Report As dialog box to save a report after customizing it.

Properties

You can customize and save a report by specifying properties such as the name and description.

Report Name

Displays the name of the report. The original report name is displayed by default. You can modify the report name as per the customization. Report name cannot exceed 255 characters.

Description

Specifies the description of the customization made on the report. Description cannot exceed 150 characters.

Command buttons

The command buttons enable you to perform the following tasks:

Save

Saves the customized report.

Cancel

Cancels the recent changes and closes the Save Customized Report As dialog box.

Save Custom Report dialog box

You can use the Save Custom Report dialog box to save a custom report after making additional changes to the custom report.

Properties

You can save a custom report by specifying properties such as the description.

Report Name

Displays the name of the custom report. This field cannot be edited.

Description

Specifies the description of the customization made on the custom report. Description cannot exceed 150 characters .

Command buttons

The command buttons enable you to perform the following tasks:

Save

Saves the custom report.

Cancel

Cancels the recent changes and closes the Save Custom Report dialog box.

Import Report dialog box

You can use the Import Report dialog box to import reports from .rptdesign files.

Properties

You can import a report by specifying the report file name, report name, and report description.

Select Report File

Enables you to select the .rptdesign file that you want to import.



In Google Chrome, the fakepath of the .rptdesign file is displayed. In Mozilla Firefox, only the .rptdesign file name is displayed. In Internet Explorer, the complete path of the .rptdesign file is displayed.

Name

Displays the name of the report. This field is empty by default. You can enter a name for the imported report.

Description

Specifies the description of the imported report. The description cannot exceed 150 characters.

· Select database user with report schema role

Select, or create, a database user if you are importing reports from the Storage Automation Store.

Command buttons

The command buttons enable you to perform the following tasks:

Import

Validates the selected .rptdesign file, and imports the report.

Cancel

Cancels the import operation, and closes the Import Report dialog box.

Configuring backup and restore operations

You can create backups of Unified Manager and use the restore feature to restore the

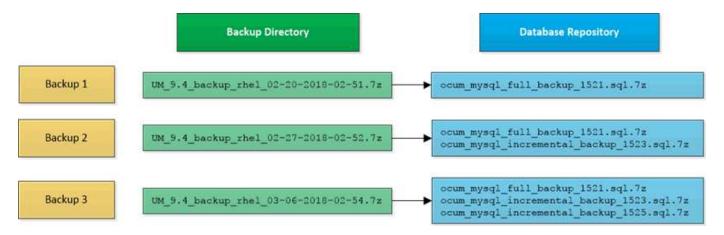
backup to the same (local) system or a new (remote) system in case of a system failure or data loss.

What a database backup is

A backup is a copy of the Unified Manager database and configuration files that you can use in case of a system failure or data loss. You can schedule a backup to be written to a local destination or to a remote destination. It is highly recommended that you define a remote location that is external to the Unified Manager host system.

A backup consists of a single file in the backup directory and one or more files in the database repository directory. The file in the backup directory is very small because it contains only a pointer to the files located in the database repository directory that are required to recreate the backup.

The first time you generate a backup a single file is created in the backup directory and a full backup file is created in the database repository directory. The next time you generate a backup a single file is created in the backup directory and an incremental backup file is created in the database repository directory that contains the differences from the full backup file. This process continues as you create additional backups, up to the maximum retention setting, as shown in the following figure.





Do not rename or remove any of the backup files in these two directories or any subsequent restore operation will fail.

If you write your backup files to the local system, you should initiate a process to copy the backup files to a remote location so they will be available in case you have a system issue that requires a complete restore.

Before beginning a backup operation, Unified Manager performs an integrity check to verify that all the required backup files and backup directories exist and are writable. It also checks that there is enough space on the system to create the backup file.

Note that you can restore a backup only on the same version of Unified Manager. For example, if you created a backup on Unified Manager 9.4, the backup can be restored only on Unified Manager 9.4 systems.

Configuring database backup settings

You can configure the Unified Manager database backup settings to set the database backup path, retention count, and backup schedules. You can enable daily or weekly scheduled backups. By default, scheduled backups are disabled.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have a minimum of 150 GB of space available in the location you define as the backup path.

It is recommended that you use a remote location that is external to the Unified Manager host system.

- When Unified Manager is installed on a Linux system, verify that the "jboss" user has write permissions to the backup directory.
- You should not schedule backup operations to occur immediately after a new cluster has been added while Unified Manager is collecting 15 days of historical performance data.

About this task

More time is required the first time a backup is performed than for subsequent backups because the first backup is a full backup. A full backup can be over 1 GB and can take three to four hours. Subsequent backups are incremental and require less time.

Steps

- 1. In the toolbar, click , and then click Management > Database Backup.
- 2. In the Management/Database Backup page, click Actions > Database Backup Settings.
- 3. Configure the appropriate values for a backup path and retention count.

The default value for retention count is 10; you can use 0 for creating unlimited backups.

- In the Schedule Frequency section, select the Enable checkbox, and then specify a daily or weekly schedule.
 - Daily

If you select this option, you must enter a time in 24-hour format for creating the backup. For example, if you specify 18:30, then a backup is created daily at 6:30 PM.

Weekly

If you select this option, you must specify the time and day for creating the backup. For example, if you specify the day as Monday and time as 16:30, then a weekly backup is created every Monday at 4:30 PM.

5. Click Save and Close.

What a database restore is

Database restore is the process of restoring an existing Unified Manager backup file to the same or a different Unified Manager server. You perform the restore operation from the Unified Manager console.

If you are performing a restore operation on the same (local) system, and the backup files are all stored locally, you can run the restore command using the default location. If you are performing a restore operation on a different Unified Manager system (a remote system), you must copy the backup file, or files, from secondary storage to the local disk before running the restore command.

During the restore process, you are logged out of Unified Manager. You can log in to the system after the restore process is complete.

The restore feature is version-specific and platform-specific. You can restore a Unified Manager backup only on the same version of Unified Manager. Unified Manager supports backup and restore in the following platform scenarios:

- · Virtual appliance to virtual appliance
- Virtual appliance to Red Hat Enterprise Linux or CentOS
- Red Hat Enterprise Linux to Red Hat Enterprise Linux or CentOS
- · Windows to Windows

If you are restoring the backup image to a new server, after the restore operation completes you need to generate a new HTTPS security certificate and restart the Unified Manager server. You will also need to reconfigure SAML authentication settings, if they are required, when restoring the backup image to a new server.



Old backup files cannot be used to restore an image after Unified Manager has been upgraded to a newer version of software. To save space, all old backup files, except the newest file, are removed automatically when you upgrade Unified Manager.

Virtual appliance backup and restore process overview

The backup and restore model for Unified Manager when installed on a virtual appliance is to capture and restore an image of the full virtual application.

Because the Unified Manager backup operation on the virtual appliance does not provide a way to move the backup file off of the vApp, the following tasks enable you to complete a backup of the virtual appliance:

- 1. Power off the VM and take a VMware snapshot of the Unified Manager virtual appliance.
- 2. Make a NetApp Snapshot copy on the datastore to capture the VMware snapshot.

If the datastore is not hosted on a system running ONTAP software, follow the storage vendor guidelines to create a backup of the VMware snapshot.

- 3. Replicate the NetApp Snapshot copy, or snapshot equivalent, to alternate storage.
- 4. Delete the VMware snapshot.

You should implement a backup schedule using these tasks to ensure that the Unified Manager virtual appliance is protected if issues arise.

To restore the VM, you can use the VMware snapshot you created to restore the VM to the backup point-in-time state.

Restoring a database backup on a virtual machine

In case of data loss or data corruption, you can use the restore feature to restore Unified Manager to the previous stable state with minimal loss. You can restore the Unified Manager database on a virtual machine by using the Unified Manager maintenance console.

Before you begin

- You must have the maintenance user credentials.
- The Unified Manager backup files must be on the local system.
- The backup files must be of .7z type.

About this task

Backup compatibility is platform and version dependent. You can restore a backup from a virtual appliance to another virtual appliance, or from a virtual appliance to a Red Hat Enterprise Linux or CentOS system.



When performing a restore operation on a different virtual appliance than the system from which the original backup file was created, the maintenance user name and password on the new vApp must be the same as the credentials from the original vApp.

Steps

- 1. In the vSphere client, locate the Unified Manager virtual machine, and then select the **Console** tab.
- Click in the console window, and then log in to the maintenance console using your user name and password.
- 3. In the Main Menu, enter the number for the System Configuration option.
- In the System Configuration Menu, enter the number for the Restore from an OCUM Backup option.
- 5. When prompted, enter the absolute path of the backup file.

```
Bundle to restore from: opt/netapp/data/ocum-backup/UM_9.4.N151112.0947_backup_unix_02-25-2018-11-41.7z
```

After the restore operation is complete, you can log in to Unified Manager.

After you finish

After you restore the backup, if the OnCommand Workflow Automation server does not work, perform the following steps:

- 1. On the Workflow Automation server, change the IP address of the Unified Manager server to point to the latest machine.
- 2. On the Unified Manager server, reset the database password if the acquisition fails in step 1.

Restoring a database backup on a Linux system

If data loss or data corruption occurs, you can restore Unified Manager to the previous stable state with minimum loss of data. You can restore the Unified Manager database to a local or remote Red Hat Enterprise Linux or CentOS system.

Before you begin

· You must have Unified Manager installed on a server.

- You must have the root user credentials for the Linux host on which Unified Manager is installed.
- You must have copied the Unified Manager backup file and the contents of the database repository directory to the system on which you will perform the restore operation.

It is recommended that you copy the backup file to the default directory /data/ocum-backup. The database repository files must be copied to the /database-dumps-repo subdirectory under the /ocum-backup directory.

• The backup files must be of .7z type.

About this task

The restore feature is platform-specific and version-specific. You can restore a Unified Manager backup only on the same version of Unified Manager. You can restore a Linux backup file or a virtual appliance backup file to a Red Hat Enterprise Linux or CentOS system.



If the backup folder name contains a space, you must include the absolute path or relative path in double quotation marks.

Steps

- 1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
- 2. Log in as the root user to the host on which Unified Manager is installed.
- If Unified Manager is installed in VCS setup, then stop the Unified Manager ocie and ocieau services using Veritas Operations Manager.

After you finish

After the restore operation is complete, you can log in to Unified Manager.

Restoring a database backup on Windows

In case of data loss or data corruption, you can use the restore feature to restore Unified Manager to the previous stable state with minimal loss. You can restore the Unified Manager database to a local Windows system or a remote Windows system by using the restore command.

Before you begin

- · You must have Unified Manager installed on a server.
- You must have Windows administrator privileges.

• You must have copied the Unified Manager backup file and the contents of the database repository directory to the system on which you will perform the restore operation.

It is recommended that you copy the backup file to the default directory \ProgramData\NetApp\OnCommandAppData\ocum\backup. The database repository files must be copied to the \database dumps repo subdirectory under the \backup directory.

• The backup files must be of .7z type.

About this task

The restore feature is platform-specific and version-specific. You can restore a Unified Manager backup only on the same version of Unified Manager, and a Windows backup can be restored only on a Windows platform.



If the folder names contain a space, you must include the absolute path or relative path of the backup file in double quotation marks.

Steps

- 1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
- 2. Log in to the Unified Manager console as an administrator: um cli login -u maint username
- 3. At the command prompt, restore the backup: um backup restore -f
 <backup_file_path>/<backup_file_name>

After you finish

After the restore operation is complete, you can log in to Unified Manager.

Description of backup windows and dialog boxes

You can view the list of backups from the backup page in Unified Manager. You can view the backup name, size, and creation time for the backups listed in this page. You can modify the database backup settings from the Database Backup Settings page.

Management/Database Backup page

The Management/Database Backup page displays a list of backups created by Unified Manager and provides information about the backup name, size, creation time, and schedule.

You must have the OnCommand Administrator or Storage Administrator role.

Command buttons

Actions

Displays the Database Backup Settings dialog box, which enables you to specify a backup path, retention count, and backup schedule.

List View

The list view displays, in tabular format, information about the backups created by Unified Manager. You can use the column filters to customize the data that is displayed.

Name

Displays the name of the selected backup.

Size

Displays the size of the selected backup.

Creation Time

Displays the creation date and time of the selected backup.

Schedule

Displays the status of the backup operation. Also indicates whether it is a scheduled backup or not.

Database Backup Settings dialog box

You can use the Database Backup Settings dialog box to specify a backup path and retention count and to enable a backup schedule for a selected backup instance.

You can change the following database backup settings:

Path

Specifies the path to the location where you store the backup files. The following table specifies the backup path format, and default locations, for different operating systems:

Host operating system	Backup path format
Virtual appliance	/opt/netapp/data/ocum-backup
Red Hat Enterprise Linux or CentOS	/data/ocum-backup
Microsoft Windows	<pre>C:\ProgramData\NetApp\OnCommandAppData \ocum\backup\</pre>

Retention Count

Specifies the maximum number of backups to be retained by Unified Manager. The default value is ten.

Schedule Frequency Enable

This option enables you to specify when to schedule a backup; you can choose daily or weekly.

Daily

Specifies the daily backup schedule with the time.

Weekly

Specifies the weekly backup schedule with the day and time.

Command buttons

· Save and Close

Saves the backup file and closes the dialog box. Unified Manager saves the backup file in the following format: um_um_version_backup_os_timestamp.7z.

Cancel

Closes the Database Backup Settings dialog box without saving your changes.

Using Unified Manager REST APIs

You can use REST APIs to help manage your clusters by viewing the health, capacity, and performance information captured by Unified Manager.

Accessing REST APIs using the Swagger API web page

REST APIs are exposed through the Swagger web page. You can access the Swagger web page to display the Unified Manager REST API documentation, as well as to manually issue an API call.

Before you begin

- You must have one of the following roles: Operator, Storage Administrator, or OnCommand Administrator.
- You must know the IP address or fully qualified domain name of the Unified Manager server on which you
 want to execute the REST APIs.

About this task

An example is provided for each REST API in the Swagger web page to help explain the objects and attributes you can use to return the information you are interested in reviewing.

Steps

1. Access the Unified Manager REST APIs.

Option	Description
From the Unified Manager web UI:	From the Menu Bar, click the Help button and then select API Documentation .
From the browser window:	Using the Unified Manager server IP address or FQDN, enter the URL to access the REST API page in the format https:// <unified_manager_ip_address_orname>/apidocs/. For example, https://10.10.10.10/apidocs/</unified_manager_ip_address_orname>

A list of API resource types, or categories, is displayed.

2. Click an API resource type to display the APIs in that resource type.

List of available REST APIs

You should be aware of the available REST APIs in Unified Manager so you can plan how you may use the APIs. The API calls are organized under the various resource types or categories.

You must refer to the Swagger web page for a complete list of the available API calls, as well as the details of each call.

The management API calls are organized according to the following categories:

- Aggregates
- Clusters
- Events
- LIFs
- LUNs
- Namespaces
- Nodes
- Ports
- SVMs
- Volumes

When you select one of the categories a list appears that shows the API sub-category along with a versioned sub-category, for example:

- · /aggregates
- /v1/aggregates

The newest version of the REST APIs are listed without a version number in the URL. You should always use the newest version of the API to integrate with Unified Manager.

Managing and monitoring Infinite Volumes

You can monitor the capacity and availability of your storage virtual machines (SVMs) with Infinite Volume. You can manage the content placement in your storage virtual machine (SVM) with Infinite Volume by creating rules and data policy.

Viewing the details of SVMs with Infinite Volume

You can use the Health/Storage Virtual Machines inventory page to view detailed information about storage virtual machines (SVMs) with Infinite Volume that are monitored by Unified Manager. You can view details such as the capacity, configuration, and data policy and rules associated with the Infinite Volume.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume must be added to the Unified Manager database.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the **Health/Storage Virtual Machines** inventory page, use the column filter in **Allowed Volume Type** to list the Infinite Volumes that are monitored.
- 3. View the complete details of the SVM with Infinite Volume by clicking the SVM name.

Viewing the constituents of an Infinite Volume

You can use the Health/Volumes inventory page to view the list of constituents in your Infinite Volume. You can view details such as the constituent state, the SVM with Infinite Volume that contains the constituent, junction path of the constituent, aggregate that contains the constituent, as well as the available, used, and total data capacity of the constituent.

Before you begin

The following requirements must be met:

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume must be added to the Unified Manager database.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. Click the name of a SVM with Infinite Volume.
- 3. In the Health/Storage Virtual Machine details page, click Volumes in the right Related Devices pane.

The list of constituents is displayed in the Health/Volumes inventory page.

Editing the Infinite Volume threshold settings

When you need to address any issues in your Infinite Volume's storage space, you can edit the threshold settings of the Infinite Volume's capacity based on your organization's requirements. When a threshold is crossed, events are generated, and you receive notifications if you have configured alerts for such events.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select a SVM with Infinite Volume.
- 3. In the Health/Storage Virtual Machine details page, click Actions > Edit Thresholds.
- 4. In the Edit SVM with Infinite Volume Thresholds dialog box, modify the thresholds as required.
- 5. Click Save and Close.

Editing the threshold settings of storage classes

When you need to address any issues related to storage space in your storage classes, you can edit the threshold settings of the storage class capacity based on your organization's requirements. When the threshold is crossed, events are generated, and you receive notifications if you have configured alerts for such events.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the **Health/Storage Virtual Machines** inventory page, select an SVM with Infinite Volume.
- 3. In the Health/Storage Virtual Machine details page, click Actions > Edit Thresholds.
- 4. In the Edit Storage Class Thresholds dialog box, modify the thresholds as required.
- 5. Click Save and Close.

Understanding Infinite Volumes

An Infinite Volume is a logical storage unit that you can use to provide a large, scalable data container with a single namespace and a single mount point. Understanding some of the basic concepts of Infinite Volumes helps you to monitor and manage your SVMs with Infinite Volume.

What an Infinite Volume is

An Infinite Volume is a single, scalable volume that can store up to 2 billion files and tens of petabytes of data.

With an Infinite Volume, you can manage multiple petabytes of data in one large logical entity and clients can retrieve multiple petabytes of data from a single junction path for the entire volume.

An Infinite Volume uses storage from multiple aggregates on multiple nodes. You can start with a small Infinite Volume and expand it nondisruptively by adding more disks to its aggregates or by providing it with more aggregates to use.

Maximum number of files an Infinite Volume can store

In most cases, an Infinite Volume can hold up to 2 billion files. If an Infinite Volume is relatively small, its maximum number of files might be less than 2 billion.

The maximum number of files that an Infinite Volume can hold is determined by the size of its namespace constituent. If the namespace constituent is 10 TB, the Infinite Volume can hold 2 billion files. If the namespace constituent is less than 10 TB, the Infinite Volume can hold proportionally fewer files.

The size of the namespace constituent is roughly proportional to the size of the Infinite Volume, depending on several factors, such as the namespace constituent's 10 TB maximum size, the available space in the aggregate that holds the namespace constituent, and the SnapDiff setting.

For a two-node Infinite Volume or a multi-node Infinite Volume without SnapDiff enabled, setting the Infinite Volume to a size of 80 TB or greater typically creates a namespace constituent of 10 TB.

The file count not only includes regular files, but also other file system structures, such as directories and symbolic links.

What a storage class is

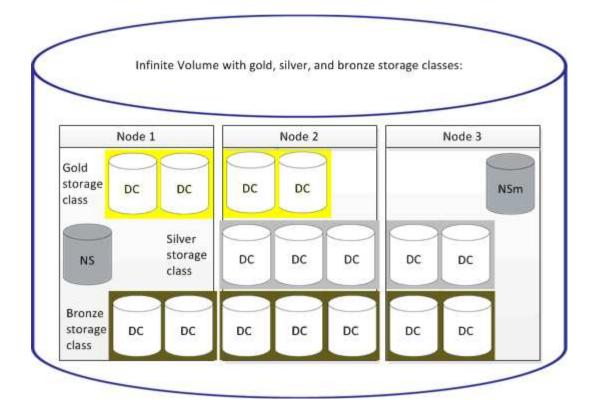
A storage class is a definition of aggregate characteristics and volume settings. You can define different storage classes and associate one or more storage classes with an Infinite Volume. You must use OnCommand Workflow Automation to define workflows for your storage class requirements and to assign storage classes to Infinite Volumes.

You can define the following characteristics for a storage class:

- Aggregate characteristics, such as the type of disks to use
- · Volume settings, such as compression, deduplication, and volume guarantee

For example, you can define a storage class that uses only aggregates with SAS disks and the following volume settings: thin provisioning with compression and deduplication enabled.

The following diagram illustrates an Infinite Volume that spans multiple nodes and uses the following storage classes: gold, silver, and bronze. Each storage class can span two or more nodes within an Infinite Volume. The diagram also illustrates the placement of data constituents in each storage class.



What a namespace constituent is

Each Infinite Volume has a single namespace constituent that maps directory information and file names to the file's physical data location within the Infinite Volume.

Clients are not aware of the namespace constituent and do not interact directly with it. The namespace constituent is an internal component of the Infinite Volume.

What data constituents are

In an Infinite Volume, data is stored in multiple separate data constituents. Data constituents store only the data from a file, not the file's name.

Clients are not aware of data constituents. When a client requests a file from an Infinite Volume, the node retrieves the file's data from a data constituent and returns the file to the client.

Each Infinite Volume typically has dozens of data constituents. For example, a 6 PB Infinite Volume that contains 1 billion files might have 60 data constituents located on aggregates from 6 nodes.

What a namespace mirror constituent is

A namespace mirror constituent is an intracluster data protection mirror copy of the namespace constituent in an Infinite Volume. The namespace mirror constituent performs two roles: It provides data protection of the namespace constituent, and it supports SnapDiff for incremental tape backup of Infinite Volumes.

Creating rules

You can add new rules to your data policy to determine the placement of data that is

written to the Infinite Volume. You can create rules either by using rule templates that are defined in Unified Manager or create custom rules.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

Creating rules using templates

You can add new rules by using rule templates defined by Unified Manager to determine the placement of data that is written to the SVM with Infinite Volume. You can create rules based on file types, directory paths, or owners.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- 3. Click the **Data Policy** tab.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

- 4. Click Create.
- 5. In the Create Rule dialog box, choose an appropriate rule template from the drop-down list.

The template is based on three categories: file type, owner, or directory path.

- 6. Based on the template selected, add the necessary conditions in the Matching Criteria area.
- 7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.
- 8. Click Create.

The new rule you created is displayed in the Data Policy tab.

- 9. Preview any other changes made to the data policy.
- 10. Click **Activate** to activate the changes in the rule properties in the SVM.

Creating custom rules

Based on your data center requirements, you can create custom rules and add them to a data policy to determine the placement of data that is written to the SVM with Infinite Volume. You can create custom rules from the Create Rule dialog box without using any existing template.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

- In the left navigation pane, click Health > SVMs.
- 2. In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- 3. Click Data Policy.
- 4. Click Create.
- 5. In the Create Rule dialog box, select Custom rule from the Template list.
- 6. In the Matching Criteria area, add conditions as required.

Conditions enable you to create a rule based on file types, directory paths, or owners. A combination of these conditions are the condition sets. For example, you can have a rule: "Place all .mp3 owned by John in bronze storage class."

- 7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.
- 8. Click Create.

The newly created rule is displayed in the Data Policy tab.

- 9. Preview any other changes made to the data policy.
- Click Activate to activate the changes in the rule properties in the SVM.

Viewing rules

You can view the list of rules you created from the Data Policy tab before modifying the data policy for your SVM with Infinite Volume.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- 3. Click Data Policy.

Results

The list of rules in the data policy for the selected SVM is displayed. You can use Filter by Storage Class to view rules about a specific storage class.

Editing template-based rules

You can edit a rule that was created using the rule templates from the Edit Rule dialog box. You can add, modify, or delete rule properties such as directory paths, file types, and owners. You can also modify the rule name and the storage class associated with the rule.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select an appropriate SVM.
- Click Data Policy.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

- 4. Select the rule for which you want to include new conditions or condition sets.
- 5. Click Edit.
- 6. In the **Edit Rule** dialog box, edit the rule as required:

If you want to	Do this
Add a new rule property	Click Add.
Delete a rule property	Click Delete by selecting the appropriate rule property.
Modify a rule property	Double-click the appropriate rule property, and then modify as required.

Click Update.

- 8. Verify that your modifications are applied to the rule by expanding the rule in the **Data Policy** tab.
- 9. Preview any other changes made to the data policy.
- 10. Click **Activate** to activate the changes to the rule properties in the SVM.

Editing custom rules

You can edit a rule to include new conditions or condition sets in the rule. For example, if you want to include new directory paths along with the owner names, you can do so from the Edit Rule dialog box.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select an appropriate SVM.
- 3. Click Data Policy.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

- 4. Select the rule for which you want to include new conditions or condition sets.
- 5. Click Edit.
- 6. In the **Edit Rule** dialog box, add new conditions or condition sets:

If you want to add	Click
A new condition	The icon.
A new condition set	Add Condition Set.

- 7. Click Update.
- 8. Verify that your modifications are applied to the rule by expanding the rule in the **Data Policy** tab.
- 9. Preview any other changes made to the data policy.
- 10. Click **Activate** to activate the changes in the rule properties in the SVM.

Deleting rules

You can delete a rule from a data policy when it is no longer required. For example, you might want to delete a rule on a particular directory that is no longer valid.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select a appropriate SVM.
- 3. Click Data Policy.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

4. Select the rule that you want to delete, and then click **Delete**.



You cannot delete the default rule.

- 5. Preview any other changes made to the data policy.
- 6. Click **Activate** to activate the changes in the rule properties in the SVM.

Previewing changes to your data policy

You should preview any changes that you have made to your rules in a data policy before you submit the changes in the data policy to the SVM with Infinite Volume for activation.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

- In the left navigation pane, click Health > SVMs.
- 2. In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- 3. Click Data Policy.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

4. Modify the data policy as required.

Data policy modifications can include creating new rules, editing existing rules, deleting existing rules, or reordering the rules.

- Click Activate.
- 6. In the Summary of Changes to Data Policy Configuration window, preview the changes to your data

policy, and then click Activate to activate the changes in the data policy in the SVM with Infinite Volume.

Exporting a data policy configuration

You can export a data policy configuration from Unified Manager to a file. For example, after you have taken the required backup, and in the event of a disaster, you can export the data policy configuration from the primary.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab, which is used while performing this task, is displayed only for SVMs with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- 3. Click Data Policy.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

- 4. Click Export.
- In the browser-specific dialog box, specify the location to which the data policy configuration has to be exported.

Results

The data policy configuration is exported as a JSON file in the specified location.

Importing a data policy configuration

You can import a data policy configuration from a file, modify the data policy, and then activate the changes in the SVM with Infinite Volume. For example, in the event of a disaster, you can import an already defined data policy to the secondary and modify the policy as required.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you import a data policy configuration, your existing rules are overwritten.

The Data Policy tab is displayed only for SVMs with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- 3. Click **Data Policy**.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

- 4. Click Import.
- 5. In the **Import Data Policy** dialog box, specify the data policy that you want to import by providing the absolute path of the data policy file.
- 6. Click Import.
- 7. Click **Activate** to activate the imported rules in the SVM.

Understanding rules and data policy

Understanding the concepts about rules and data policy help you to manage your Infinite Volumes efficiently.

What rules and data policies are

A *rule* determines the placement of files (data) in a storage virtual machine (SVM) with Infinite Volume. A collection of such rules is known as a *data policy*.

Rule

Rules mainly consist of a set of predefined conditions and information that determine where to place files in the Infinite Volume. When a file is placed in the Infinite Volume, the attributes of that file are matched with the list of rules. If attributes match the rules, then that rule's placement information determines the storage class where the file is placed. A default rule in the data policy is used to determine the placement of files if the attributes do not match any of the rules in the rule list.

For example, if you have a rule, "Place all files of type .mp3 in the bronze storage class.", all .mp3 files that are written to the Infinite Volume would be placed in the bronze storage class.

· Data policy

A data policy is a list of rules. Each SVM with Infinite Volume has its own data policy. Each file that is added to the Infinite Volume is compared to its data policy's rules to determine where to place that file. The data policy enables you to filter incoming files based on the file attributes and place these files in the appropriate storage classes.

What the default rule is

The default rule is the rule present in the data policy of a storage virtual machine (SVM) with Infinite Volume. It is used to determine the placement of data written to the Infinite Volume when none of the conditions in the existing rules match with the data being written.

The default rule is always the last rule in a data policy and cannot be reordered. For example, consider a data

policy with three rules. Rule-1 places all .pdf files in the *high_performance*storage class. Rule-2 places all files owned by the administrator and file names that end with *.xls in the *archival_constituent* storage class. The third rule is the default rule with the *low performance* storage class.

When a set of *.jpg files that are not owned by the administrator is written to the Infinite Volume, the default rule is used to place these .jpg files in the *low_performance* storage class. Rule-1 and Rule-2 are not used because the data that is written does not match these rules.

How a data policy filters data written to an Infinite Volume

A data policy automatically filters data written to the Infinite Volume into different storage classes. All files are written to the single file system in the Infinite Volume's namespace, and rules in the data policy determine which storage class stores the data for the files.

A default data policy is automatically created for a storage virtual machine (SVM) with Infinite Volume when you create the Infinite Volume. The data policy is active and contains a default rule. The default rule stores incoming data for files as follows for Infinite Volumes with and without storage classes:

For an Infinite Volume	The default data policy does this
Without storage classes	Places all incoming data for files in the Infinite Volume
With one storage class	Places all incoming data for files into the storage class
With one or more storage classes	Places all incoming data for files into the first storage class that is created

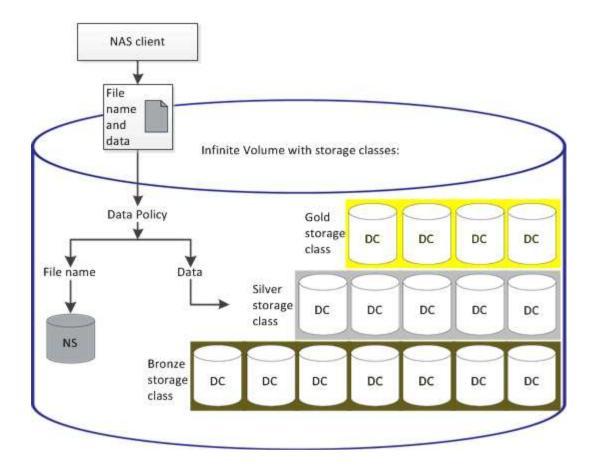


For an Infinite Volume with two or more storage classes, you should modify the data policy as soon as possible to create rules that filter data for different types of files into the different storage classes. You should modify the data policy by using Unified Manager.

The data policy does not affect the location of the files in the file system in the Infinite Volume's namespace, and storage classes are transparent to client applications. The file system in the namespace contains the file names. The data policy affects only which storage class is used to store the data for the files. Data policies are useful when you assign two or more storage classes to an Infinite Volume.

You can modify the data policy to create additional rules, but you cannot delete the data policy or its default rule.

The following diagram illustrates how a data policy filters data for an Infinite Volume. The file name is stored in the namespace constituent, and rules in the data policy specify that data for this particular file is stored in the silver storage class.



What a rule template is

A rule template is a predefined template that can be used to create rules in a data policy. A rule template enables you to create a rule based on three categories: owner, file type, and directory path.

Example of a rule template for file types

The rule template "Place all files with the specified extensions in a suitable storage class" places all the .mp3 files that are written to the Infinite Volume in a storage class that you specify.

What conditions and condition sets are

Conditions are a set of matching criteria based on rule properties—such as the file name, directory path, and owner—that define a rule. A collection of such conditions is known as a condition set. You can use conditions and condition sets only for custom rules to determine where to place content that is written into your Infinite Volume.

Conditions

For a custom rule, you can specify conditions based on rule properties such as the file name, directory path, or owner, or a combination of all the rule properties. The logic is similar to a Boolean AND operation. For example, by using conditions, you can create a custom rule to place files with .mp3 extensions and files owned by John in the directory path starting with /NS/.

Condition sets

The logic used for condition sets is similar to a Boolean OR operation. For example, by using conditions and condition sets, you can create a complex custom rule that matches either of the following conditions:

condition-1

All files owned by Mary and are placed in /NS/Eng/

condition-2

All files that have names ending with .pdfand owned by Mary

Description of Infinite Volume windows and dialog boxes

You can monitor SVMs with Infinite Volume from the respective Health/Storage Virtual Machine details page. You can manage rules and data policies from the Create Rule dialog box. You can also modify the storage class thresholds from the Edit Storage Class Thresholds dialog box.

Create Rule dialog box

You can use the Create Rule dialog box to create new rules for your data policy. For example, when you want to specify the placement of content of a certain file type, you can use the Create Rule dialog box to create the rule for your data policy.

Rule Name

Specifies the name of the new rule.

Templates area

Displays the list of rule templates. You can select an appropriate rule template from the list to create a rule for the data policy.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected. For example, if you select "Place all files with the specified owner names in a suitable storage class", **List of Owner that...** is displayed in Matching Criteria.

Add

Enables you to add a new rule property based on the rule template selected. For example, if you selected the rule template, "Place all files with the specified owner names in a suitable storage class", the **Add** button enables you to add the owner's name.

Delete

Enables you to delete a selected rule property.

Content Placement

Enables you to select an appropriate storage class for your rule from the list.

Command buttons

Create

Creates a new rule for the data policy and closes the Create Rule dialog box.

Cancel

Cancels the recent changes made to the rule and closes the Create Rule dialog box.

Edit Rule dialog box

You can use the Edit Rule dialog box to edit the properties of a rule, such as the file types, directory paths, or owners. You can also select an appropriate storage class for the rule. For example, when a certain file path is no longer valid, you can delete the file path from the corresponding rule.

Rule Name

Displays the name of the rule.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected.

Add

Enables you to add a new rule property, a new file type, a file path, or a new owner. For example, if you had specified the rule template, "Place all files with the specified owner names in a suitable storage class", the Add button enables you to add the owner's name.

Delete

Enables you to delete a selected rule property.

Content Placement area

Displays the list of storage classes. You can select an appropriate storage class for the selected rule.

Command buttons

Update

Updates the changes made to the rule and closes the Edit Rule dialog box.

Cancel

Cancels the recent changes made to the rule and closes the Edit Rule dialog box.

Edit Rule dialog box (Advanced edit)

You can use the Edit Rule dialog box to edit the properties of a rule that is not created by using a template. The rule properties you can edit include the file types, directory paths, matching criteria, or owners. You can select an appropriate storage class for the rule. For example, you can edit the conditions specified in the matching criteria of a rule.

Rule Name

Displays the name of the rule.

Matching Criteria

Displays a list of conditions related to the selected rule template. The condition list changes based on the rule template selected. You can expand the rules and modify the rule properties, as required.

Content Placement area

Displays the list of storage classes. You can select an appropriate storage class for the selected rule.

Command buttons

Update

Updates the changes made to the rule and closes the Edit Rule dialog box.

Cancel

Cancels the recent changes made to the rule and closes the Edit Rule dialog box.

Edit SVM with Infinite Volume Thresholds dialog box

You can use the Edit SVM with Infinite Volume Thresholds dialog box to modify the default threshold values of each SVM with Infinite Volume, based on your organization's requirements. The default threshold values indicate the level of activity that must be reached on the SVM before an event is triggered.

Capacity

The Capacity area enables you to set capacity threshold conditions for the selected SVM with Infinite Volume:

Space Nearly Full

Specifies the percentage at which the SVM with Infinite Volume is considered to be nearly full. It also displays the corresponding space (in GB, MB, or TB) in the Infinite Volume. For example, if you have an Infinite Volume of size 10 GB, and the Space Nearly Full threshold is 80%, then the following information is displayed: (8 GB of 10 GB).

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which the SVM with Infinite Volume is considered full. It also displays the

corresponding space (in GB, MB, or TB) in the Infinite Volume. For example, if you have an Infinite Volume of size 10 GB, and the Space Full threshold is 90%, then the following information is displayed: (9 GB of 10 GB).

You can also use the slider to set the threshold value.

Snapshot Usage Limit

Specifies the limit, in percentage, of space reserved for Snapshot copies in the Infinite Volume.

Command buttons

The command buttons enable you to perform the following tasks:

Restore to Global Defaults

Enables you to restore the threshold settings to the current values that are set at the global level.

Save

Saves all the threshold settings.

· Save and Close

Saves all the threshold settings and then closes the Edit SVM with Infinite Volume Thresholds dialog box.

Cancel

Ignores any changes to the threshold settings and closes the Edit SVM with Infinite Volume Thresholds dialog box.

Edit Storage Class Thresholds dialog box

You can use the Edit Storage Class Thresholds dialog box to modify the default threshold values of various storage classes in each SVM with Infinite Volume based on your organization's requirements. The default threshold values indicate the level of activity that must be reached on a storage class before an event is triggered.

You must have the OnCommand Administrator or Storage Administrator role.

Capacity

The Capacity area enables you to set capacity threshold conditions for the selected storage class.

Space Nearly Full

Specifies the percentage at which a storage class in the SVM with Infinite Volume is considered to be nearly full. It also displays the corresponding space (in GB, MB, or TB) in the storage class. For example, if you have a storage class of size 10 GB and the Space Nearly Full threshold is 80%, then the following information is displayed: (8 GB of 10 GB).

You can also use the slider to set the threshold value.

Space Full

Specifies the percentage at which the storage class in the SVM with Infinite Volume is considered full. It also displays the corresponding space (in GB, MB, or TB) in the storage class. For example, if you have a storage class of size 10 GB and the Space Full threshold is 90%, then the following information is displayed: (9 GB of 10GB).

You can also use the slider to set the threshold value.

Snapshot Usage Limit

Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.

Command buttons

The command buttons enable you to perform tasks for a selected volume.

· Restore to Global Defaults

Enables you to restore the threshold settings to the current values that are set at the global level.

Save

Saves all the threshold settings.

Save and Close

Saves all the threshold settings and then closes the Edit Storage Class Thresholds dialog box.

Cancel

Cancels changes (if any) to the threshold settings and closes the Edit Storage Class Thresholds dialog box.

Managing clusters

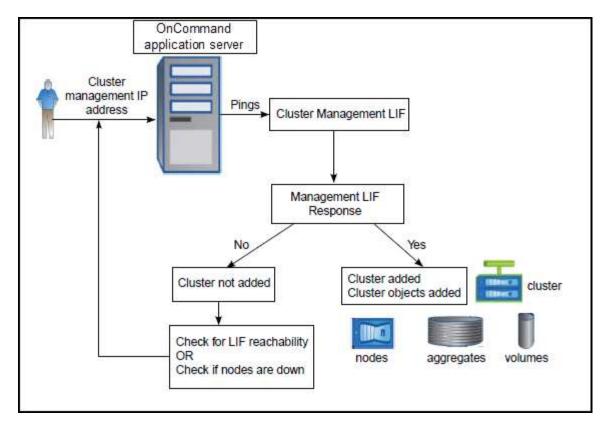
You can manage the ONTAP clusters by using Unified Manager to monitor, add, edit, and remove clusters.

How the cluster discovery process works

After you have added a cluster to Unified Manager, the server discovers the cluster objects and adds them to its database. Understanding how the discovery process works helps you to manage your organization's clusters and their objects.

The monitoring interval for collecting cluster configuration information is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster objects in the Unified Manager UI. This time frame is also true when making changes to a cluster. For example, if you add two new volumes to an SVM in a cluster, you see those new objects in the UI after the next polling interval, which could be up to 15 minutes.

The following image illustrates the discovery process:



After all the objects for a new cluster are discovered, Unified Manager starts to gather historical performance data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.



Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time.

Viewing the list of monitored clusters

You can use the Configuration/Cluster Data Sources page to view your inventory of clusters. You can view details about the clusters, such as their name or IP address and communication status.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

The list of clusters is sorted by the collection state severity level column. You can click a column header to sort the clusters by different columns.

Steps

1. In the left navigation pane, click **Configuration > Cluster Data Sources**.

Adding clusters

You can add a cluster to OnCommand Unified Manager so that you can monitor the cluster. This includes the ability to obtain cluster information such as the health, capacity, performance, and configuration of the cluster so that you can find and resolve any issues that might occur.

Before you begin

- You must have the OnCommand Administrator role or the Storage Administrator role.
- You must have the host name or cluster management IP address (IPv4 or IPv6) for the cluster.

When using the host name, it must resolve to the cluster management IP address for the cluster management LIF. If you use a node management LIF, the operation fails.

• You must have the user name and password to access the cluster.

This account must have the admin role with Application access set to ontapi, ssh, and http.

- You must know the type of protocol (HTTP or HTTPS) that is to be configured on the cluster and the port number use to connect to the cluster.
- You must have adequate space on the Unified Manager server. You are prevented from adding a cluster to the server when greater than 90% of space is already consumed.



You can add clusters which are behind a NAT/firewall by using the Unified Manager NAT IP address. Any connected Workflow Automation or SnapProtect systems must also be behind the NAT/firewall, and SnapProtect API calls must use the NAT IP address to identify the cluster.

About this task

- Each cluster in a MetroCluster configuration must be added separately.
- A single instance of Unified Manager can support a specific number of nodes. If you need to monitor an
 environment that exceeds the supported node count, you must install an additional instance of Unified
 Manager to monitor some of the clusters.
- You can monitor a single cluster by two instances of Unified Manager provided that you have configured a second cluster-management LIF on the cluster so that each instance of Unified Manager connects through a different LIF.

Steps

- 1. In the left navigation pane, click Configuration > Cluster Data Sources.
- On the Configuration/Cluster Data Sources page, click Add.
- 3. In the Add Cluster dialog box, specify the values as required, and then click Submit.
- 4. If HTTPS is selected, perform the following steps:
 - a. In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information about the cluster.
 - b. Click Yes.

Unified Manager checks the certificate only when the cluster is added initially. Unified Manager does

not check the certificate for each API call to ONTAP.

If the certificate has expired, you cannot add a new cluster. You must first renew the SSL certificate and then add the cluster.

Results

After all the objects for a new cluster are discovered (about 15 minutes), Unified Manager starts to gather historical performance data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.



Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time. Additionally, if you restart Unified Manager during the data continuity collection period, the collection will be halted and you will see gaps in the performance charts for the missing timeframe.

If you receive an error message that you cannot add the cluster, check to see if the following issues exist:



- If the clocks on the two systems are not synchronized and the Unified Manager HTTPS
 certificate start date is later than the date on the cluster. You must ensure that the clocks are
 synchronized using NTP or a similar service.
- If the cluster has reached the maximum number of EMS notification destinations the Unified Manager address cannot be added. By default only 20 EMS notification destinations can be defined on the cluster.

Editing clusters

You can modify the settings of an existing cluster, such as the host name or IP address, user name, password, protocol, and port, by using the Edit Cluster dialog box.

Before you begin

You must have the OnCommand Administrator role or the Storage Administrator role.

About this task



If you change the IP address of a cluster to an IP address of an existing monitored cluster, all data for the existing cluster is lost when the former cluster is discovered. An error message is not displayed to warn you.

Steps

- 1. In the left navigation pane, click Configuration > Cluster Data Sources.
- On the Configuration/Cluster Data Sources page, select the cluster you want to edit, and then click Edit.
- 3. In the **Edit Cluster** dialog box, modify the values as required.
- 4. Click Submit.

Removing clusters

You can remove a cluster from Unified Manager by using the Configuration/Cluster Data Sources page. For example, you can remove a cluster if cluster discovery fails, or when you want to decommission a storage system.

Before you begin

You must have the OnCommand Administrator role or the Storage Administrator role.

About this task

This task removes the selected cluster from Unified Manager. After a cluster is removed, it is no longer monitored. The instance of Unified Manager registered with the removed cluster is also unregistered from the cluster.

Removing a cluster also deletes all its storage objects, historical data, storage services, and all associated events from Unified Manager. These changes are reflected on the inventory pages and the details pages after the next data collection cycle.

Steps

- 1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
- On the Configuration/Cluster Data Sources page, select the cluster that you want to remove and click Remove.
- 3. In the Remove Data Source message dialog, click Remove to confirm the remove request.

Rediscovering clusters

You can manually rediscover a cluster from the Configuration/Cluster Data Sources page page in order to obtain the latest information about the health, monitoring status, and performance status of the cluster.

About this task

You can manually rediscover a cluster when you want to update the cluster—such as by increasing the size of an aggregate when there is insufficient space—and you want Unified Manager to discover the changes that you make.

When Unified Manager is paired with OnCommand Workflow Automation (WFA), the pairing triggers the reacquisition of the data cached by WFA.

Steps

- 1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
- On the Configuration/Cluster Data Sources page, click Rediscover.

Unified Manager rediscovers the selected cluster and displays the latest health and performance status.



You can obtain the monitoring status of the cluster from the right pane of the Dashboards/Cluster View page.

Page descriptions for data source management

You can view and manage your clusters, including adding, editing, rediscovering, and removing clusters, from a single page.

Configuration/Cluster Data Sources page

The Configuration/Cluster Data Sources page displays information about the clusters that Unified Manager is currently monitoring. This page enables you to add additional clusters, edit cluster settings, and remove clusters.

A message at the bottom of the page indicates how frequently Unified Manager collects performance data from clusters. The default collection interval is five minutes, but you can modify this interval through the maintenance console if you find that collections from large clusters are not completing on time.

Command buttons

Add

Opens the Add Cluster dialog box, which enables you to add clusters.

• Edit

Opens the Edit Cluster dialog box, which enables you to edit the settings of the selected cluster.

Remove

Removes the selected cluster and all the associated events and storage objects. After the cluster is removed, it is no longer monitored.



The cluster, its storage objects, and all associated events are removed, and the cluster is no longer monitored by Unified Manager. The instance of Unified Manager registered with the removed clustered is also unregistered from the cluster.

Rediscover

Forces a rediscover operation of the cluster so you can update the collection of health and performance data.

Clusters list

The Clusters list displays the properties of all the discovered clusters. You can click a column header to sort the clusters by that column.

Status

Displays the current discovery status of the data source. The status can be Failed (1), Completed (2), or In Progress (3).

Name

Displays the cluster name.

Note that the name might take fifteen minutes or more to appear after the cluster is first added.

Maintenance Mode

Enables you to specify the timeframe, or "maintenance window", when a cluster will be down for maintenance so that you do not receive a storm of alerts from the cluster while it is being maintained.

When maintenance mode is scheduled for the future this field displays "Scheduled", and you can hover your cursor over the field to display the scheduled time. When the cluster is in the maintenance window this field shows "Active".

Host Name or IP Address

Displays the host name, fully qualified domain name (FQDN), short name, or the IP address of the cluster-management LIF that is used to connect to the cluster.

Protocol

Displays the type of protocol that can be configured on the cluster: HTTP or HTTPS (for a secure connection).

If a connection is established with the cluster by using both protocols, HTTPS is chosen over HTTP. The default is HTTPS.

Port

Displays the port number of the cluster.

If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

User Name

Displays the user name that can be used to log in to the cluster.

Operation

Displays the current operation that is supported by the cluster data source.

The following operations are supported by the data source:

Discovery

Specifies the operation when the data source is being discovered.

Health Poll

Specifies the operation when the data source is successfully discovered and has started sampling data.

Deletion

Specifies the operation when the data source (cluster) is deleted from the respective storage objects list.

Operation State

Displays the state of the current operation. The state can be Failed, Completed, or In Progress.

Operation Start Time

The date and time the operation started.

Operation End Time

The date and time the operation ended.

Description

Any message related to the operation.

Add Cluster dialog box

You can add an existing cluster so that you can monitor the cluster and obtain information about its health, capacity, configuration, and performance.

You can add a cluster by specifying the following values:

Host Name or IP Address

Enables you to specify the host name (preferred) or the IP address (IPv4 or IPv6) of the clustermanagement LIF that is used to connect to the cluster. By specifying the host name, you will be able to match the name of the cluster across the web UI, rather than trying to correlate an IP address on one page to a host name on another page.

User Name

Enables you to specify a user name that can be used to log in to the cluster.

Password

Enables you to specify a password for the specified user name.

Protocol

Enables you to specify the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.

Port

Enables you to specify the port number used to connect to the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Edit Cluster dialog box

The Edit Cluster dialog box enables you to modify the connection settings of an existing cluster, including the IP address, port, and protocol.

You can edit the following fields:

Host Name or IP Address

Enables you to specify the FQDN, short name, or the IP address (IPv4 or IPv6) of the cluster-management LIF that is used to connect to the cluster.

User Name

Enables you to specify a user name that can be used to log in to the cluster.

Password

Enables you to specify a password for the specified user name.

Protocol

Enables you to specify the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.

Port

Enables you to specify the port number used to connect to the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Managing user access

You can create roles and assign capabilities to control user access to selected cluster objects. You can identify users who have the required capabilities to access selected objects within a cluster. Only these users are provided access to manage the cluster objects.

Adding users

You can add local users or database users by using the Management/Users page. You can also add remote users or groups that belong to an authentication server. You can assign roles to these users and, based on the privileges of the roles, users can manage the storage objects and data with Unified Manager, or view the data in a database.

Before you begin

- You must have the OnCommand Administrator role.
- To add a remote user or group, you must have enabled remote authentication and configured your authentication server.
- If you plan to configure SAML authentication so that an identity provider (IdP) authenticates users accessing the graphical interface, make sure these users are defined as "remote" users.

Access to the UI is not allowed for users of type "local" or "maintenance" when SAML authentication is enabled.

About this task

If you add a group from Windows Active Directory, then all direct members and nested subgroups can authenticate to Unified Manager, unless nested subgroups are disabled. If you add a group from OpenLDAP or other authentication services, then only the direct members of that group can authenticate to Unified Manager.

Steps

- 1. In the toolbar, click [6], and then click **Users** in the left Management menu.
- 2. On the Management/Users page, click Add.
- In the Add User dialog box, select the type of user that you want to add, and enter the required information.

When entering the required user information, you must specify an email address that is unique to that user. You must avoid specifying email addresses that are shared by multiple users.

Click Add.

Editing the user settings

You can edit user settings—such as the email address and role—that are specified each user. For example, you might want to change the role of a user who is a storage operator, and assign storage administrator privileges to the user.

Before you begin

You must have the OnCommand Administrator role.

About this task

When you modify the role that is assigned to a user, the changes are applied when either of the following actions occur:

- The user logs out and logs back in to Unified Manager.
- · Session timeout of 24 hours is reached.

Steps

- 1. In the toolbar, click , and then click **Users** in the left Management menu.
- 2. In the Management/Users page, select the user for which you want to edit settings, and click Edit.
- In the Edit User dialog box, edit the appropriate settings that are specified for the user.
- 4. Click Save.

Testing a remote user or remote group

You can validate that a remote user or remote group can access the Unified Manager server by using the authentication settings that are specified for your authentication servers.

Before you begin

- You must have enabled remote authentication, and configured your authentication settings so that the Unified Manager server can validate the remote user or remote group.
- You must have the OnCommand Administrator role.

Steps

- 1. In the toolbar, click , and then click **Users** in the left Management menu.
- In the Management/Users page, select a remote user or remote group that you want to validate, and then click Test.

Viewing users

You can use the Management/Users page to view the list of users who manage storage objects and data using Unified Manager. You can view details about the users, such as the user name, type of user, email address, and the role that is assigned to the users.

Before you begin

You must have the OnCommand Administrator role.

Steps

1. In the toolbar, click , and then click **Users** in the left Management menu.

The list of users is displayed in the Management/Users page.

Deleting users or groups

You can delete one or more users from the management server database to prevent specific users from accessing Unified Manager. You can also delete groups so that all the users in the group can no longer access the management server.

Before you begin

When you are deleting remote groups, you must have reassigned the events that are assigned to the users
of the remote groups.

If you are deleting local users or remote users, the events that are assigned to these users are automatically unassigned.

• You must have the OnCommand Administrator role.

Steps

- 1. In the toolbar, click 💽, and then click **Users** in the left Management menu.
- 2. In the Management/Users page, select the users or groups that you want to delete, and then click Delete.
- 3. Click Yes to confirm the deletion.

Changing the local user password

You can change your local user 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 remote users cannot be changed using these steps. To change a remote user password, contact your password administrator. To change the maintenance user password, see Using the Maintenance Console.

Steps

- 1. Log in to Unified Manager.
- 2. From the top menu bar, click the user icon and then click **Change Password**.

The **Change Password** option is not displayed if you are a remote user.

- 3. In the Change Password dialog box, enter the current password and the new password.
- 4. Click Save.

After you finish

If Unified Manager is configured in a high-availability configuration, you must change the password on the second node of the setup. Both instances must have same password.

What the maintenance user does

The maintenance user is created during the installation of Unified Manager on a Red Hat Enterprise Linux or CentOS system. The maintenance user name is the "umadmin" user. The maintenance user has the OnCommand administrator role in the web UI, and that user can create subsequent users and assign them roles.

The maintenance user, or umadmin user, can also access the Unified Manager maintenance console.

What RBAC is

RBAC (role-based access control) provides the ability to control who has access to various features and resources in the OnCommand Unified Manager server.

What role-based access control does

Role-based access control (RBAC) enables administrators to manage groups of users by defining roles. If you need to restrict access for specific functionality to selected administrators, you must set up administrator accounts for them. If you want to restrict the information that administrators can view and the operations they can perform, you must apply roles to the administrator accounts you create.

The management server uses RBAC for user login and role permissions. If you have not changed the management server's default settings for administrative user access, you do not need to log in to view them.

When you initiate an operation that requires specific privileges, the management server prompts you to log in. For example, to create administrator accounts, you must log in with Administrator account access.

Definitions of user types

A user type specifies the kind of account the user holds and includes remote users, remote groups, local users, database users, and maintenance users. Each of these types has its own role, which is assigned by a user with the role of OnCommand Administrator.

Unified Manager user types are as follows:

Maintenance user

Created during the initial configuration of Unified Manager. The maintenance user then creates additional users and assigns roles. The maintenance user is also the only user with access to the maintenance console. When Unified Manager is installed on a Red Hat Enterprise Linux or CentOS system, the maintenance user is given the user name "umadmin."

Local user

Accesses the Unified Manager UI and performs functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.

· Remote group

A group of users that access the Unified Manager UI using the credentials stored on the authentication server. The name of this account should match the name of a group stored on the authentication server. All users within the remote group are given access to the Unified Manager UI using their individual user credentials. Remote groups can perform functions according to their assigned roles.

· Remote user

Accesses the Unified Manager UI using the credentials stored on the authentication server. A remote user performs functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.

· Database user

Has read-only access to data in the Unified Manager database, has no access to the Unified Manager web interface or the maintenance console, and cannot execute API calls.

Definitions of user roles

The maintenance user or OnCommand administrator assigns a role to every user. Each role contains certain privileges. The scope of activities that you can perform in Unified Manager depends on the role you are assigned and which privileges the role contains.

Unified Manager includes the following predefined user roles:

Operator

Views storage system information and other data collected by Unified Manager, including histories and capacity trends. This 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. This role enables the storage administrator to configure thresholds and to create alerts and other storage management-specific options and policies.

OnCommand Administrator

Configures settings unrelated to storage management. This role enables the management of users, security certificates, database access, and administrative options, including authentication, SMTP, networking, and AutoSupport.



When Unified Manager is installed on Linux systems, the initial user with the OnCommand Administrator role is automatically named "umadmin".

Integration Schema

This role enables read-only access to Unified Manager database views for integrating Unified Manager with OnCommand Workflow Automation (WFA).

Report Schema

This role enables read-only access to reporting and other database views directly from the Unified Manager database. The databases that can be viewed include:

- netapp model view
- netapp_performance
- · ocum
- ocum_report
- ocum report birt
- opm
- scalemonitor

Unified Manager user roles and capabilities

Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

The following table displays the functions that each user role can perform:

Function	Operator	Storage Administrator	OnCommand Administrator	Integration Schema	Report Schema
View storage system information	•	•	•	•	•

Function	Operator	Storage Administrator	OnCommand Administrator	Integration Schema	Report Schema
View other data, such as histories and capacity trends	•	•	•	•	•
View, assign, and resolve events	•	•	•		
View storage service objects, such as SVM associations and resource pools	•	•	•		
View threshold policies	•	•	•		
Manage storage service objects, such as SVM associations and resource pools		•	•		
Define alerts		•	•		
Manage storage management options		•	•		
Manage storage management policies		•	•		
Manage users			•		
Manage administrative options			•		
Define threshold policies			•		
Manage database access			•		

Function	Operator	Storage Administrator	OnCommand Administrator	Integration Schema	Report Schema
Manage integration with WFA and provide access to the database views				•	
Provide read- only access to reporting and other database views					•
Schedule and save reports	•	•	•		
Import and delete imported reports			•		

Description of user access windows and dialog boxes

Based on the RBAC settings, you can add users from the Management/Users page and assign appropriate roles to those users to access and monitor your clusters.

Management/Users page

The Management/Users page displays a list of your users and groups, and provides information such as the name, type of user, and email address. You can also use this page to perform tasks such as adding, editing, deleting, and testing users.

Command buttons

The command buttons enable you to perform the following tasks for selected users:

Add

Displays the Add User dialog box, which enables you to add a local user, a remote user, a remote group, or a database user.

You can add remote users or groups only if your authentication server is enabled and configured.

• Edit

Displays the Edit User dialog box, which enables you to edit the settings for the selected user.

Delete

Deletes the selected users from the management server database.

Test

Enables you to validate whether a remote user or group is present in the authentication server.

You can perform this task only if your authentication server is enabled and configured.

List view

The List view displays, in tabular format, information about the users that are created. You can use the column filters to customize the data that is displayed.

Name

Displays the name of the user or group.

Type

Displays the type of user: Local User, Remote User, Remote Group, Database User, or Maintenance User.

• Email

Displays the email address of the user.

Role

Displays the type of role that is assigned to the user: Operator, Storage Administrator, OnCommand Administrator, Integration Schema, or Report Schema.

Add User dialog box

You can create local users or database users, or add remote users or remote groups, and assign roles so that these users can manage storage objects and data using Unified Manager.

You can add a user by completing the following fields:

Type

Enables you to specify the type of user you want to create.

Name

Enables you to specify a user name that a user can use to log in to Unified Manager.

Password

Enables you to specify a password for the specified user name. This field is displayed only when you are adding a local user or a database user.

· Confirm Password

Enables you to reenter your password to ensure the accuracy of what you entered in the Password field.

This field is displayed only when you are adding a local user or a database user.

Email

Enables you to specify an email address for the user; the email address specified must be unique to the user name. This field is displayed only when you are adding a remote user or a local user.

Role

Enables you to assign a role to the user and defines the scope of activities that the user can perform. The role can be OnCommand Administrator, Storage Administrator, Operator, Integration Schema, or Report Schema.

Command buttons

The command buttons enable you to perform the following tasks:

Add

Adds the user and closes the Add User dialog box.

Cancel

Cancels the changes and closes the Add User dialog box.

Edit User dialog box

The Edit User dialog box enables you to edit only certain settings, depending on the selected user.

Details

The Details area enables you to edit the following information about a selected user:

Type

This field cannot be edited.

Name

This field cannot be edited.

Password

Enables you to edit the password when the selected user is a database user.

Confirm Password

Enables you to edit the confirmed password when the selected user is a database user.

Email

Enables you to edit the email address of the selected user. This field can be edited when the selected user is a local user, LDAP user, or maintenance user.

Role

Enables you to edit the role that is assigned to the user. This field can be edited when the selected user is a local user, remote user, or remote group.

Command buttons

The command buttons enable you to perform the following tasks:

Save

Saves the changes and closes the Edit User dialog box.

Cancel

Cancels the changes and closes the Edit User dialog box.

Managing authentication

You can enable authentication using either LDAP or Active Directory on the Unified Manager server and configure it to work with your servers to authenticate remote users.

Additionally, you can enable SAML authentication so that remote users are authenticated through a secure identity provider (IdP) before they can log into the Unified Manager web UI.

Enabling remote authentication

You can enable remote authentication so that the Unified Manager server can communicate with your authentication servers. The users of the authentication server can access the Unified Manager graphical interface to manage storage objects and data.

Before you begin

You must have the OnCommand Administrator role.



The Unified Manager server must be connected directly with the authentication server. You must disable any local LDAP clients such as SSSD (System Security Services Daemon) or NSLCD (Name Service LDAP Caching Daemon).

About this task

You can enable remote authentication using either Open LDAP or Active Directory. If remote authentication is disabled, remote users cannot access Unified Manager.

Remote authentication is supported over LDAP and LDAPS (Secure LDAP). Unified Manager uses 389 as the default port for non-secure communication, and 636 as the default port for secure communication.



The certificate that is used to authenticate users must conform to the X.509 format.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, select Enable Remote Authentication.
- 3. In the **Authentication Service** field, select the type of service and configure the authentication service.

For Authentication type	Enter the following information
Active Directory	 Authentication server administrator name in one of following formats:
	° domainname \ username
	° username@domainname
	° Bind Distinguished Name (using the appropriate LDAP notation)
	Administrator password
	Base distinguished name (using the appropriate LDAP notation)
Open LDAP	Bind distinguished name (in the appropriate LDAP notation)
	Bind password
	Base distinguished name

If the authentication of an Active Directory user takes a long time or times out, the authentication server is probably taking a long time to respond. Disabling support for nested groups in Unified Manager might reduce the authentication time.

If you select the Use Secure Connection option for the authentication server, then Unified Manager communicates with the authentication server using the Secure Sockets Layer (SSL) protocol.

- 4. Add authentication servers, and test the authentication.
- 5. Click Save and Close.

Disabling nested groups from remote authentication

If you have remote authentication enabled, you can disable nested group authentication so that only individual users, and not group members, can remotely authenticate to Unified Manager. You can disable nested groups when you want to improve Active Directory authentication response time.

Before you begin

- You must have the OnCommand Administrator role.
- Disabling nested groups is only applicable when using Active Directory.

About this task

Disabling support for nested groups in Unified Manager might reduce the authentication time. If nested group support is disabled, and if a remote group is added to Unified Manager, individual users must be members of the remote group to authenticate to Unified Manager.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, check the Disable Nested Group Lookup box.
- 3. Click Save.

Setting up authentication services

Authentication services enable the authentication of remote users or remote groups in an authentication server before providing them access to Unified Manager. You can authenticate users by using predefined authentication services (such as Active Directory or OpenLDAP), or by configuring your own authentication mechanism.

Before you begin

- You must have enabled remote authentication.
- · You must have the OnCommand Administrator role.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the **Setup** options page, click **Management Server > Authentication**.
- 3. Select one of the following authentication services:

If you select	Then do this
Active Directory	a. Enter the administrator name and password.
	 Specify the base distinguished name of the authentication server.
	For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is cn=ou, dc=domain, dc=com.

If you select	Then do this
OpenLDAP	a. Enter the bind distinguished name and bind password.
	b. Specify the base distinguished name of the authentication server.
	For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is cn=ou, dc=domain, dc=com.
Others	a. Enter the bind distinguished name and bind password.
	b. Specify the base distinguished name of the authentication server.
	For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is cn=ou, dc=domain, dc=com.
	c. Specify the LDAP protocol version that is supported by the authentication server.
	d. Enter the user name, group membership, user group, and member attributes.



If you want to modify the authentication service, you must delete any existing authentication servers, and then add new authentication servers.

4. Click Save and Close.

Adding authentication servers

You can add authentication servers and enable remote authentication on the management server so that remote users within the authentication server can access Unified Manager.

Before you begin

- The following information must be available:
 - Host name or IP address of the authentication server
 - Port number of the authentication server
- You must have enabled remote authentication and configured your authentication service so that the management server can authenticate remote users or groups in the authentication server.
- You must have the OnCommand Administrator role.

About this task

If the authentication server that you are adding is part of a high-availability (HA) pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, click Management Server > Authentication.
- 3. Enable or disable the **Use secure connection authentication** option:

If you want to	Then do this
Enable it	a. In Enable remote authentication checkbox, select the Use Secure Connection option.
	b. In the Authentication Servers area, click Add .
	 c. In the Add Authentication Server dialog box, enter the authentication name or IP address (IPv4 or IPv6) of the server.
	d. In the Authorize Host dialog box, click View Certificate.
	e. In the View Certificate dialog box, verify the certificate information, and then click Close .
	f. In the Authorize Host dialog box, click Yes .
	When you enable the Use Secure Connection authentication option, Unified Manager communicates with the authentication server and displays the certificate. Unified Manager uses 636 as default port for secure communication and port number 389 for non-secure communication.
Disable it	 a. In the Enable remote authentication checkbox, clear the Use Secure Connection option.
	b. In the Authentication Servers area, click Add .
	c. In the Add Authentication Server dialog box, specify either the host name or IP address (IPv4 or IPv6) of the server, and the port details.d. Click Add.

The authentication server that you added is displayed in the Servers area.

4. Perform a test authentication to confirm that you can authenticate users in the authentication server that

Testing the configuration of authentication servers

You can validate the configuration of your authentication servers to ensure that the management server is able to communicate with them. You can validate the configuration by searching for a remote user or remote group from your authentication servers, and authenticating them using the configured settings.

Before you begin

- You must have enabled remote authentication, and configured your authentication service so that the Unified Manager server can authenticate the remote user or remote group.
- You must have added your authentication servers so that the management server can search for the remote user or remote group from these servers and authenticate them.
- You must have the OnCommand Administrator role.

About this task

If the authentication service is set to Active Directory, and if you are validating the authentication of remote users who belong to the primary group of the authentication server, information about the primary group is not displayed in the authentication results.

Steps

- 1. In the toolbar, click 💽, and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, click Test Authentication.
- 3. In the **Test User** dialog box, specify the user name and password of the remote user or the user name of the remote group, and then click **Test**.

If you are authenticating a remote group, you must not enter the password.

Editing authentication servers

You can change the port that the Unified Manager server uses to communicate with your authentication server.

Before you begin

You must have the OnCommand Administrator role.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, check the Disable Nested Group Lookup box.
- 3. In the **Authentication Servers** area, select the authentication server that you want to edit, and then click **Edit**.
- In the Edit Authentication Server dialog box, edit the port details.

Deleting authentication servers

You can delete an authentication server if you want to prevent the Unified Manager server from communicating with the authentication server. For example, if you want to change an authentication server that the management server is communicating with, you can delete the authentication server and add a new authentication server.

Before you begin

You must have the OnCommand Administrator role.

About this task

When you delete an authentication server, remote users or groups of the authentication server will no longer be able to access Unified Manager.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the **Setup/Authentication** page, select one or more authentication servers that you want to delete, and then click **Delete**.
- 3. Click **Yes** to confirm the delete request.

If the **Use Secure Connection** option is enabled, then the certificates associated with the authentication server are deleted along with the authentication server.

Authentication with Active Directory or OpenLDAP

You can enable remote authentication on the management server and configure the management server to communicate with your authentication servers so that users within the authentication servers can access You can enable remote authentication on the management server and configure the management server to communicate with your authentication servers so that users within the authentication servers can access Unified Manager.

You can use one of the following predefined authentication services or specify your own authentication service:

· Microsoft Active Directory



You cannot use Microsoft Lightweight Directory Services.

OpenLDAP

You can select the required authentication service and add the appropriate authentication servers to enable the remote users in the authentication server to access Unified Manager. The credentials for remote users or groups are maintained by the authentication server. The management server uses the Lightweight Directory Access Protocol (LDAP) to authenticate remote users within the configured authentication server.

For local users who are created in Unified Manager, the management server maintains its own database of user names and passwords. The management server performs the authentication and does not use Active Directory or OpenLDAP for authentication.

Enabling SAML authentication

You can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Before you begin

- · You must have configured remote authentication and verified that it is successful.
- You must have created at least one Remote User, or a Remote Group, with the OnCommand Administrator role.
- The Identity provider (IdP) must be supported by Unified Manager and it must be configured.
- · You must have the IdP URL and metadata.
- · You must have access to the IdP server.

About this task

After you have enabled SAML authentication from Unified Manager, users cannot access the graphical user interface until the IdP has been configured with the Unified Manager server host information. So you must be prepared to complete both parts of the connection before starting the configuration process. The IdP can be configured before or after configuring Unified Manager.

Only remote users will have access to the Unified Manager graphical user interface after SAML authentication has been enabled. Local users and Maintenance users will not be able to access the UI. This configuration does not impact users who access the maintenance console, the Unified Manager commands, or ZAPIs.



Unified Manager is restarted automatically after you complete the SAML configuration on this page.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- In the Setup/Authentication page, select the SAML Authentication tab.
- 3. Select the **Enable SAML authentication** checkbox.

The fields required to configure the IdP connection are displayed.

4. Enter the IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP server.

If the IdP server is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URI to populate the IdP Metadata field automatically.

5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.

You can configure the IdP server with this information at this time.

Click Save.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

7. Click Confirm and Logout and Unified Manager is restarted.

Results

The next time authorized remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the IdP login page instead of the Unified Manager login page.

After you finish

If not already completed, access your IdP and enter the Unified Manager server URI and metadata to complete the configuration.



When using ADFS as your identity provider, the Unified Manager GUI does not honor the ADFS timeout and will continue to work until the Unified Manager session timeout is reached. When Unified Manager is deployed on Windows, Red Hat, or CentOS, you can change the GUI session timeout using the following Unified Manager CLI command: um option set absolute.session.timeout=00:15:00This command sets the Unified Manager GUI session timeout to 15 minutes.

Identity provider requirements

When configuring Unified Manager to use an identity provider (IdP) to perform SAML authentication for all remote users, you need to be aware of some required configuration settings so that the connection to Unified Manager is successful.

You must enter the Unified Manager URI and metadata into the IdP server. You can copy this information from the Unified ManagerSAML Authentication page. Unified Manager is considered the service provider (SP) in the Security Assertion Markup Language (SAML) standard.

Supported encryption standards

- Advanced Encryption Standard (AES): AES-128 and AES-256
- Secure Hash Algorithm (SHA): SHA-1 and SHA-256

Validated identity providers

- · Shibboleth
- Active Directory Federation Services (ADFS)

ADFS configuration requirements

 You must define three claim rules in the following order that are required for Unified Manager to parse ADFS SAML responses for this relying party trust entry.

Claim rule	Value
SAM-account-name	Name ID
SAM-account-name	urn:oid:0.9.2342.19200300.100.1.1
Token groups — Unqualified Name	urn:oid:1.3.6.1.4.1.5923.1.5.1.1

- You must set the authentication method to "Forms Authentication" or users may receive an error when logging out of Unified Manager when using Internet Explorer. Follow these steps:
 - a. Open the ADFS Management Console.
 - b. Click on the Authentication Policies folder on the left tree view.
 - c. Under Actions on the right, click Edit Global Primary Authentication Policy.
 - d. Set the Intranet Authentication Method to "Forms Authentication" instead of the default "Windows Authentication".
- In some cases login through the IdP is rejected when the Unified Manager security certificate is CA-signed. There are two workarounds to resolve this issue:
 - Follow the instructions identified in the link to disable the revocation check on the ADFS server for chained CA cert associated relying party:

http://www.torivar.com/2016/03/22/adfs-3-0-disable-revocation-check-windows-2012-r2/

Have the CA server reside within the ADFS server to sign the Unified Manager server cert request.

Other configuration requirements

- The Unified Manager clock skew is set to 5 minutes, so the time difference between the IdP server and the Unified Manager server cannot be more than 5 minutes or authentication will fail.
- When users attempt to access Unified Manager using Internet Explorer they might see the message The
 website cannot display the page. If this occurs, make sure these users uncheck the option for "Show
 friendly HTTP error messages" in Tools > Internet Options > Advanced.

Changing the identity provider used for SAML authentication

You can change the identity provider (IdP) that Unified Manager uses to authenticate remote users.

Before you begin

- You must have the IdP URL and metadata.
- · You must have access to the IdP.

About this task

The new IdP can be configured before or after configuring Unified Manager.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, select the SAML Authentication tab.
- Enter the new IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP.

If the IdP is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URL to populate the IdP Metadata field automatically.

- 4. Copy the Unified Manager metadata URI, or save the metadata to an XML text file.
- 5. Click Save Configuration.

A message box displays to confirm that you want to change the configuration.

6. Click OK.

After you finish

Access the new IdP and enter the Unified Manager server URI and metadata to complete the configuration.

The next time authorized remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the new IdP login page instead of the old IdP login page.

Disabling SAML authentication

You can disable SAML authentication when you want to stop authenticating remote users through a secure identity provider (IdP) before they can log into the Unified Manager web UI. When SAML authentication is disabled, the configured directory service providers, such as Active Directory or LDAP, perform sign-on authentication.

About this task

After you disable SAML authentication, Local users and Maintenance users will be able to access the graphical user interface in addition to configured Remote users.

You can also disable SAML authentication using the Unified Manager maintenance console if you do not have access to the graphical user interface.



Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- In the Setup/Authentication page, select the SAML Authentication tab.
- 3. Uncheck the Enable SAML authentication checkbox.
- 4. Click Save.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

5. Click Confirm and Logout and Unified Manager is restarted.

Results

The next time remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the Unified Manager login page instead of the IdP login page.

After you finish

Access your IdP and delete the Unified Manager server URI and metadata.

Description of authentication windows and dialog boxes

You can enable LDAP authentication from the Setup/Authentication page.

Setup/Authentication page

You can use the Setup/Authentication page to configure Unified Manager to authenticate remote users who attempt to log into the Unified Manager web UI.

Using the Remote Authentication page you can configure Unified Manager to communicate with your authentication server to authenticate remote users.

Using the SAML Authentication page you can configure Unified Manager to communicate with a secure Identity provider (IdP) to authenticate remote users.

Remote Authentication page

You can use the Remote Authentication page to configure Unified Manager to communicate with your authentication server to authenticate remote users who attempt to log into the Unified Manager web UI.

You must have the OnCommand Administrator or Storage Administrator role.

After you select the Enable remote authentication checkbox, you can enable remote authentication using an authentication server.

Authentication Service

Enables you to configure the management server to authenticate users in directory service providers, such as Active Directory, OpenLDAP, or specify your own authentication mechanism. You can specify an authentication service only if you have enabled remote authentication.

Active Directory

Administrator Name

Specifies the administrator name of the authentication server.

Password

Specifies the password to access the authentication server.

Base Distinguished Name

Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is cn=ou, dc=domain, dc=com.

Disable Nested Group Lookup

Specifies whether to enable or disable the nested group lookup option. By default this option is disabled. If you use Active Directory, you can speed up authentication by disabling support for nested groups.

Use Secure Connection

Specifies the authentication service used for communicating with authentication servers.

OpenLDAP

Bind Distinguished Name

Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server.

Bind Password

Specifies the password to access the authentication server.

Base Distinguished Name

Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is cn=ou, dc=domain, dc=com.

Use Secure Connection

Specifies that Secure LDAP is used for communicating with LDAPS authentication servers.

Others

Bind Distinguished Name

Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server that you have configured.

Bind Password

Specifies the password to access the authentication server.

Base Distinguished Name

Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is cn=ou, dc=domain, dc=com.

Protocol Version

Specifies the Lightweight Directory Access Protocol (LDAP) version that is supported by your authentication server. You can specify whether the protocol version must be automatically detected or set the version to 2 or 3.

User Name Attribute

Specifies the name of the attribute in the authentication server that contains user login names to be authenticated by the management server.

Group Membership Attribute

Specifies a value that assigns the management server group membership to remote users based on an attribute and value specified in the user's authentication server.

UGID

If the remote users are included as members of a GroupOfUniqueNames object in the authentication server, this option enables you to assign the management server group membership to the remote users based on a specified attribute in that GroupOfUniqueNames object.

Disable Nested Group Lookup

Specifies whether to enable or disable the nested group lookup option. By default this option is disabled. If you use Active Directory, you can speed up authentication by disabling support for nested groups.

Member

Specifies the attribute name that your authentication server uses to store information about the individual members of a group.

User Object Class

Specifies the object class of a user in the remote authentication server.

Group Object Class

Specifies the object class of all groups in the remote authentication server.

Use Secure Connection

Specifies the authentication service used for communicating with authentication servers.



If you want to modify the authentication service, ensure that you delete any existing authentication servers and add new authentication servers.

Authentication Servers area

The Authentication Servers area displays the authentication servers that the management server communicates with to find and authenticate remote users. The credentials for remote users or groups are maintained by the authentication server.

Command buttons

Enables you to add, edit, or delete authentication servers.

Add

Enables you to add an authentication server.

If the authentication server that you are adding is part of an high-availability pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.

Edit

Enables you to edit the settings for a selected authentication server.

Delete

Deletes the selected authentication servers.

Name or IP Address

Displays the host name or IP address of the authentication server that is used to authenticate the user on the management server.

Port

Displays the port number of the authentication server.

Test Authentication

This button validates the configuration of your authentication server by authenticating a remote user or group.

While testing, if you specify only the user name, the management server searches for the remote user in the authentication server, but does not authenticate the user. If you specify both the user name and password, the management server searches and authenticates the remote user.

You cannot test the authentication if remote authentication is disabled.

SAML Authentication page

You can use the SAML Authentication page to configure Unified Manager to authenticate remote users using SAML though a secure identity provider (IdP) before they can to log into the Unified Manager web UI.

- You must have the OnCommand Administrator role to create or modify the SAML configuration.
- · You must have configured remote authentication.
- You must have configured at least one remote user or remote group.

After remote authentication and remote users have been configured, you can select the Enable SAML authentication checkbox to enable authentication using a secure identity provider.

• IdP URI

The URI to access the IdP from the Unified Manager server. Example URIs are listed below.

ADFS example URI:

https://win2016-dc.ntap2016.local/federationmetadata/2007-06/federationmetadata.xml

Shibboleth example URI:

https://centos7.ntap2016.local/idp/shibboleth

IdP Metadata

The IdP metadata in XML format.

If the IdP URL is accessible from the Unified Manager server, you can click the **Fetch IdP Metadata** button to populate this field.

Host System (FQDN)

The FQDN of the Unified Manager host system as defined during installation. You can change this value if necessary.

Host URI

The URI to access the Unified Manager host system from the IdP.

Host Metadata

The host system metadata in XML format.

Managing security certificates

You can configure HTTPS in the Unified Manager server to monitor and manage your clusters over a secure connection.

Viewing the HTTPS security certificate

You can compare the HTTPS certificate details to the retrieved certificate in your browser to ensure that your browser's encrypted connection to Unified Manager is not being intercepted.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

Viewing the certificate enables you to verify the content of a regenerated certificate, or to view alternate URL names from which you can access Unified Manager.

Steps

1. In the toolbar, click , and then click HTTPS Certificate from the Setup menu.

The HTTPS certificate is displayed at the top of the page

After you finish

If you need to view more detailed information about the security certificate than what is displayed on the HTTPS Certificate page, you can view the connection certificate in your browser.

Generating an HTTPS security certificate

You might generate a new HTTPS security certificate for multiple reasons, including if you want to sign with a different Certificate Authority or if the current security certificate has expired. The new certificate replaces the existing certificate.

Before you begin

You must have the OnCommand Administrator role.

About this task

If you do not have access to the Unified Manager web UI, you can regenerate the HTTPS certificate with the same values using the maintenance console.

Steps

- 1. In the toolbar, click , and then click HTTPS Certificate from the Setup menu.
- 2. Click Regenerate HTTPS Certificate.

The Regenerate HTTPS Certificate dialog box is displayed.

3. Select one of the following options depending on how you want to generate the certificate:

If you want to	Do this
Regenerate the certificate with the current values	Click the Regenerate Using Current Certificate Attributes option.

If you want to	Do this	
Generate the certificate using different values	Click the *Update the Current Certificate Attributes* option. The Common Name and Alternative Names fields will use the values from the existing certificate if you do not enter new values. The other fields do not require values, but you can enter values, for example, for the City, State, and Country if you want those values to be populated in the certificate.	
	i	You can select the "Exclude local identifying information (e.g. localhost)" checkbox if you want to remove the local identifying information from the Alternative Names field in the certificate. When this checkbox is selected only what you enter in the field is used in the Alternative Names field. When left blank the resulting certificate will not have an Alternative Names field at all.

- 4. Click **Yes** to regenerate the certificate.
- 5. Restart the Unified Manager server so that the new certificate takes effect.

After you finish

Verify the new certificate information by viewing the HTTPS certificate.

Restarting the Unified Manager virtual machine

You can restart the virtual machine from the maintenance console of Unified Manager. You must restart after generating a new security certificate or if there is a problem with the virtual machine.

Before you begin

The virtual appliance is powered on.

You are logged in to the maintenance console as the maintenance user.

About this task

You can also restart the virtual machine from vSphere by using the **Restart Guest** option. See the VMware documentation for more information.

Steps

- 1. Access the maintenance console.
- 2. Select System Configuration > Reboot Virtual Machine.

Downloading an HTTPS certificate signing request

You can download a certification request for the current HTTPS security certificate so that you can provide the file to a Certificate Authority to sign. A CA-signed certificate helps prevent man-in-the-middle attacks and provides better security protection than a self-signed certificate.

Before you begin

You must have the OnCommand Administrator role.

Steps

- 1. In the toolbar, click , and then click HTTPS Certificate from the Setup menu.
- 2. Click Download HTTPS Certificate Signing Request.
- 3. Save the <hostname>.csr file.

After you finish

You can provide the file to a Certificate Authority to sign, and then install the signed certificate.

Installing an HTTPS security certificate

You can upload and install a security certificate after a Certificate Authority has signed and returned it. The file that you upload and install must be a signed version of the existing self-signed certificate. A CA-signed certificate helps prevent man-in-the middle attacks and provides better security protection than a self-signed certificate.

Before you begin

You must have completed the following actions:

- Downloaded the Certificate Signing Request file and had it signed by a Certificate Authority
- Saved the certificate chain in PEM format
- Included all certificates in the chain, from the Unified Manager server certificate to the root signing certificate, including any intermediate certificates present

You must have the OnCommand Administrator role.

Steps

- 1. In the toolbar, click , and then click HTTPS Certificate from the Setup menu.
- Click Install HTTPS Certificate.
- 3. In the dialog box that is displayed, click Choose file... to locate the file to upload.

4. Select the file, and then click Install to install the file.

Example certificate chain

The following example shows how the certificate chain file might appear:

```
----BEGIN CERTIFICATE----

<*Server certificate*>
----END CERTIFICATE----

----BEGIN CERTIFICATE----

<*Intermediate certificate \#1 \(if present\)*>
----END CERTIFICATE----

----BEGIN CERTIFICATE----

<*Intermediate certificate \#2 \(if present\)*>
----END CERTIFICATE----

<*Root signing certificate*>
----END CERTIFICATE----
```

Page descriptions for certificate management

You can use the HTTPS Certificate page to view the current security certificates and to generate new HTTPS certificates.

HTTPS Certificate page

The HTTPS Certificate page enables you to view the current security certificate, download a certificate signing request, generate a new HTTPS certificate, or install a new HTTPS certificate.

If you have not generated a new HTTPS certificate, the certificate that appears on this page is the certificate that was generated during installation.

Command buttons

The command buttons enable you to perform the following operations:

Download HTTPS Certificate Signing Request

Downloads a certification request for the currently installed HTTPS certificate. Your browser prompts you to save the <nostname>.csr file so that you can provide the file to a Certificate Authority to sign.

Install HTTPS Certificate

Enables you to upload and install a security certificate after a Certificate Authority has signed and returned it. The new certificate is in effect after you restart the management server.

Regenerate HTTPS Certificate

Enables you to generate an HTTPS certificate, which replaces the current security certificate. The new

certificate is in effect after you restart Unified Manager.

Regenerate HTTPS Certificate dialog box

The Regenerate HTTPS Certificate dialog box enables you to customize the security information and then generate a new HTTPS certificate with that information.

The current certificate information appears on this page.

The "Regenerate Using Current Certificate Attributes" and "Update the Current Certificate Attributes" selection enables you to regenerate the certificate with the current information or generate a certificate with new information.

Common Name

Required. The fully qualified domain name (FQDN) that you wish to secure.

In Unified Manager high availability configurations, use the virtual IP address.

• Email

Optional. An email address to contact your organization; typically the email address of the certificate administrator or IT department.

Company

Optional. Typically the incorporated name of your company.

Department

Optional. The name of the department in your company.

City

Optional. The city location of your company.

State

Optional. The state or province location, not abbreviated, of your company.

Country

Optional. The country location of your company. This is typically a two-letter ISO code of the country.

Alternative Names

Required. Additional, non-primary domain names that can be used to access this server in addition to the existing localhost or other network addresses. Separate each alternate name with a comma.

Select the "Exclude local identifying information (e.g. localhost)" checkbox if you want to remove the local identifying information from the Alternative Names field in the certificate. When this checkbox is selected only what you enter in the field is used in the Alternative Names field. When left blank the resulting certificate will not have an Alternative Names field at all.

Troubleshooting

Troubleshooting information helps you to identify and resolve issues you encounter when using Unified Manager.

Changing the Unified Manager host name

At some point, you might want to change the host name of the system on which you have installed Unified Manager. For example, you might want to rename the host to more easily identify your Unified Manager servers by type, workgroup, or monitored cluster group.

The steps required to change the host name are different depending on whether Unified Manager is running on a VMware ESXi server, on a Red Hat or CentOS Linux server, or on a Microsoft Windows server.

Changing the Unified Manager virtual appliance host name

The network host is assigned a name when the Unified Manager virtual appliance is first deployed. You can change the host name after deployment. If you change the host name, you must also regenerate the HTTPS certificate.

Before you begin

You must be logged in to Unified Manager as the maintenance user, or have the OnCommand Administrator role assigned to you to perform these tasks.

About this task

You can use the host name (or the host IP address) to access the Unified Manager web UI. If you configured a static IP address for your network during deployment, then you would have designated a name for the network host. If you configured the network using DHCP, the host name should be taken from the DNS. If DHCP or DNS is not properly configured, the host name "OnCommand" is automatically assigned and associated with the security certificate.

Regardless of how the host name was assigned, if you change the host name, and intend to use the new host name to access the Unified Manager web UI, you must generate a new security certificate.

If you access the web UI by using the server's IP address instead of the host name, you do not have to generate a new certificate if you change the host name. However, it is the best practice to update the certificate so that the host name in the certificate matches the actual host name.

If you change the host name in Unified Manager, you must manually update the host name in OnCommand Workflow Automation (WFA). The host name is not updated automatically in WFA.

The new certificate does not take effect until the Unified Manager virtual machine is restarted.

Steps

1. Generate an HTTPS security certificate

If you want to use the new host name to access the Unified Manager web UI, you must regenerate the HTTPS certificate to associate it with the new host name.

2. Restart the Unified Manager virtual machine

After you regenerate the HTTPS certificate, you must restart the Unified Manager virtual machine.

Changing the Unified Manager host name on Linux systems

At some point, you might want to change the host name of the Red Hat Enterprise Linux or CentOS machine on which you have installed Unified Manager. For example, you might want to rename the host to more easily identify your Unified Manager servers by type, workgroup, or monitored cluster group when you list your Linux machines.

Before you begin

You must have root user access to the Linux system on which Unified Manager is installed.

About this task

You can use the host name (or the host IP address) to access the Unified Manager web UI. If you configured a static IP address for your network during deployment, then you would have designated a name for the network host. If you configured the network using DHCP, the host name should be taken from the DNS server.

Regardless of how the host name was assigned, if you change the host name and intend to use the new host name to access the Unified Manager web UI, you must generate a new security certificate.

If you access the web UI by using the server's IP address instead of the host name, you do not have to generate a new certificate if you change the host name. However, it is the best practice to update the certificate, so that the host name in the certificate matches the actual host name. The new certificate does not take effect until the Linux machine is restarted.

If you change the host name in Unified Manager, you must manually update the host name in OnCommand Workflow Automation (WFA). The host name is not updated automatically in WFA.

Steps

- 1. Log in as the root user to the Unified Manager system that you want to modify.
- 2. Stop the Unified Manager software and the associated MySQL software by entering the following commands in the order shown: service ocieau stopservice ocie stopservice mysqld stop
- Change the host name using the Linux hostnamectl command: hostnamectl set-hostname new_FQDN

hostnamectl set-hostname nuhost.corp.widget.com

- 4. Regenerate the HTTPS certificate for the server:/opt/netapp/essentials/bin/cert.sh create
- 5. Restart the network service: service network restart
- 6. After the service is restarted, verify whether the new host name is able to ping itself: ping new hostname

ping nuhost

This command should return the same IP address that was set earlier for the original host name.

7. After you complete and verify your host name change, restart Unified Manager by entering the following commands in the order shown: service mysqld startservice ocie startservice ocieau start

Adding disk space to the Unified Manager database directory

The Unified Manager database directory contains all of the health and performance data collected from ONTAP systems. Some circumstances may require that you increase the size of the database directory.

For example, the database directory may get full if Unified Manager is collecting data from a large number of clusters where each cluster has many nodes. You will receive a warning event when the database directory is 90% full, and a critical event when the directory is 95% full.



No additional data is collected from clusters after the directory reaches 95% full.

The steps required to add capacity to the data directory are different depending on whether Unified Manager is running on a VMware ESXi server, on a Red Hat or CentOS Linux server, or on a Microsoft Windows server.

Adding space to the data disk of the VMware virtual machine

If you need to increase the amount of space on the data disk for the Unified Manager database, you can add capacity after installation by increasing disk space on disk 3.

Before you begin

- You must have access to the vSphere Client.
- · The virtual machine must have no snapshots stored locally.
- · You must have the maintenance user credentials.

About this task

We recommend that you back up your virtual machine before increasing the size of virtual disks.

Steps

- 1. In the vSphere client, select the Unified Manager virtual machine, and then add more disk capacity to data disk 3. See the VMware documentation for details.
- 2. In the vSphere client, select the Unified Manager virtual machine, and then select the **Console** tab.
- 3. Click in the console window, and then log in to the maintenance console using your user name and password.
- 4. In the **Main Menu**, enter the number for the **System Configuration** option.
- 5. In the **System Configuration Menu**, enter the number for the **Increase Data Disk Size** option.

Adding space to the data directory of the Linux host

If you allotted insufficient disk space to the <code>/opt/netapp/data</code> directory to support Unified Manager when you originally set up the Linux host and then installed Unified Manager, you can add disk space after installation by increasing disk space on the

/opt/netapp/data directory.

Before you begin

You must have root user access to the Red Hat Enterprise Linux or CentOS Linux machine on which Unified Manager is installed.

About this task

We recommend that you back up the Unified Manager database before increasing the size of the data directory.

Steps

- 1. Log in as root user to the Linux machine on which you want to add disk space.
- 2. Stop the Unified Manager service and the associated MySQL software in the order shown: service ocieau stopservice ocie stopservice mysqld stop
- 3. Create a temporary backup folder (for example, /backup-data) with sufficient disk space to contain the data in the current /opt/netapp/data directory.
- 4. Copy the content and privilege configuration of the existing /opt/netapp/data directory to the backup data directory: cp -rp /opt/netapp/data/* /backup-data
- 5. If SE Linux is enabled:
 - a. Get the SE Linux type for folders on existing /opt/netapp/data folder:

```
se_type= ls -Z /opt/netapp/data | awk '{print $4}'| awk -F: '{print $3}'|
head -1
```

The system returns a confirmation similar to the following:

```
echo $se_type
mysqld_db_t
```

- b. Run the chcon command to set the SE Linux type for the backup directory: chcon -R
 --type=mysqld_db_t /backup-data
- 6. Remove the contents of the /opt/netapp/data directory:

```
a. cd /opt/netapp/data
```

- b. rm -rf *
- 7. Expand the size of the /opt/netapp/data directory to a minimum of 750 GB through LVM commands or by adding extra disks.



Mounting the /opt/netapp/data directory on an NFS export or CIFS share is not supported.

8. Confirm that the /opt/netapp/data directory owner (mysql) and group (root) are unchanged: ls -ltr / | grep opt/netapp/data

The system returns a confirmation similar to the following:

```
drwxr-xr-x. 17 mysql root 4096 Aug 28 13:08 data
```

9. If SE Linux is enabled, confirm that the context for the <code>/opt/netapp/data</code> directory is still set to <code>mysqld_db_t</code>: touch <code>/opt/netapp/data/abc</code>`ls <code>-Z</code> <code>/opt/netapp/data/abc</code>

The system returns a confirmation similar to the following:

```
-rw-r--r-. root root unconfined_u:object_r:mysqld_db_t:s0 /opt/netapp/data/abc
```

- 10. Copy the contents from backup-data, back to the expanded /opt/netapp/data directory: cp -rp /backup-data/* /opt/netapp/data/
- 11. Start the MySQL service: service mysqld start
- 12. After the MySQL service is started, start the ocie and ocieau services in the order shown: service ocie start``service ocieau start
- 13. After all of the services are started, delete the backup folder /backup-data: rm -rf /backup-data

Adding space to the logical drive of the Microsoft Windows server

If you need to increase the amount of disk space for the Unified Manager database, you can add capacity to the logical drive on which Unified Manager is installed.

Before you begin

You must have Windows administrator privileges.

About this task

We recommend that you back up the Unified Manager database before adding disk space.

Steps

- 1. Log in as administrator to the Windows server on which you want to add disk space.
- 2. Follow the step that corresponds to method you want to use to add more space:

Option	Description	
On a physical server, add capacity to the logical drive on which the Unified Manager server is installed.	Follow the steps in the Microsoft topic: Extend a Basic Volume	
On a physical server, add a hard disk drive.	Follow the steps in the Microsoft topic: Adding Hard Disk Drives	

Option	Description
On a virtual machine, increase the size of a disk partition.	Follow the steps in the VMware topic: Increasing the size of a disk partition

Changing the performance statistics collection interval

The default collection interval for performance statistics is 5 minutes. You can change this interval to 10 or 15 minutes if you find that collections from large clusters are not finishing within the default time. This setting affects the collection of statistics from all clusters that this instance of Unified Manager is monitoring.

Before you begin

You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.

About this task

The issue of performance statistics collections not finishing on time is indicated by the banner messages Unable to consistently collect from cluster <cluster_name> or Data collection is taking too long on cluster <cluster name>.

You should change the collection interval only when required because of a statistics collections issue. Do not change this setting for any other reason.



Changing this value from the default setting of 5 minutes can affect the number and frequency of performance events that Unified Manager reports. For example, system-defined performance thresholds trigger events when the policy is exceeded for 30 minutes. When using 5-minute collections, the policy must be exceeded for six consecutive collections. For 15-minute collections the policy must be exceeded for only two collection periods.

A message at the bottom of the Cluster Data Sources page indicates the current statistical data collection interval.

Steps

1. Log in using SSH as the maintenance user to the Unified Manager host.

The Unified Managermaintenance console prompts are displayed.

- 2. Type the number of the menu option labeled **Performance Polling Interval Configuration**, and then press Enter.
- 3. If prompted, enter the maintenance user password again.
- 4. Type the number for the new polling interval that you want to set, and then press Enter.

After you finish

If you changed the Unified Manager collection interval to 10 or 15 minutes, and you have a current connection

to an external data provider (such as Graphite), you must change the data provider transmit interval so that it is equal to, or greater, than the Unified Manager collection interval.

Enabling periodic AutoSupport

You can choose to have specific, predefined messages sent automatically to technical support to ensure correct operation of your environment, and to assist you in maintaining the integrity of your environment.

Before you begin

You must be logged in as the maintenance user.

About this task

You must activate AutoSupport in order to receive the benefits of NetApp Active IQ.

Steps

- 1. In the toolbar, click , and then click **AutoSupport** from the **Setup** menu.
- 2. Select the **Enable Periodic AutoSupport** check box.
- 3. If required, define the name, port, and authentication information for the HTTP proxy server.
- 4. Click Save.

Sending on-demand AutoSupport messages

You can send Unified Manager system information to technical support for assistance with troubleshooting issues. The AutoSupport message contains diagnostic system information and detailed data about the Unified Manager server.

Before you begin

You must be logged in as the maintenance user.

Steps

- 1. In the toolbar, click , and then click **AutoSupport** from the **Setup** menu.
- 2. Perform one or both of the following actions:

If you want to send the AutoSupport message to	Do this
Technical support	Select the Send to Technical Support check box.
A specific email recipient	Select the Send to Email Recipient check box, and enter the email address of the recipient.

3. If required, define the name, port, and authentication information for the HTTP proxy server, and click **Save**.

4. Click Generate and send AutoSupport.

Setup/AutoSupport page

The Setup/AutoSupport page enables you to view the AutoSupport description, enable periodic AutoSupport, or send an on-demand AutoSupport message.

Information area

System ID

Displays the system ID for this Unified Manager server.

Periodic AutoSupport area

Enables you to have specific, predefined messages to technical support for issue diagnosis and resolution periodically generated.

Enable Periodic AutoSupport

Indicates that you want to enable the periodic AutoSupport functionality.

On-Demand AutoSupport area

You can generate and send an on-demand message to technical support, a specified email recipient, or both:

Send to Technical Support

Indicates that you want to send an on-demand message to technical support for any issues that have occurred.

Send to Email Recipient

Indicates that you want to send an on-demand message to a specified recipient for any issues that have occurred.

Generate and Send AutoSupport

Generates and sends an on-demand message to technical support, a specified email recipient, or both for any issues that have occurred.

HTTP Proxy area

You can designate a proxy to provide Internet access in order to send AutoSupport content to support if your environment does not provide direct access from the Unified Manager server.

Use HTTP proxy

Check this box to identify the server being used as the HTTP proxy.

Enter the host name or IP address of the proxy server, and the port number used to connect to the server.

· Use authentication

Check this box if you need to provide authentication information to access the server being used as the HTTP proxy.

Enter the user name and the password required to authenticate with the HTTP proxy.



HTTP proxies that provide only Basic Authentication are not supported.

Unknown authentication error

Issue

When you are performing an authentication-related operation such as adding, editing, deleting, or testing remote users or groups, the following error message might be displayed: Unknown authentication error.

Cause

This problem can occur if you have set an incorrect value for the following options:

- Administrator Name of the Active Directory authentication service
- Bind Distinguished Name of the OpenLDAP authentication service

Corrective action

- a. In the toolbar, click [5], and then click **Setup > Authentication**.
- b. Based on the authentication service that you have selected, enter the appropriate information for Administrator Name or Bind Distinguished Name.
- c. Click **Test Authentication** to test the authentication with the details that you specified.
- d. Click Save and Close.

User not found

Issue

When you are performing an authentication-related operation such as adding, editing, deleting, or testing remote users or groups, the following error message is displayed: User not found.

Cause

This problem can occur if the user exists in the AD server or LDAP server, and if you have set the base distinguished name to an incorrect value.

Corrective action

- a. In the toolbar, click , and then click **Setup > Authentication**.
- b. Enter the appropriate information for base distinguished name.
- c. Click Save and Close.

Issue with adding LDAP using Other authentication services

Issue

When you select Others as the Authentication service, the user and groupObjectClass retain the values from the previously selected template. If the LDAP server does not use the same values, the operation might fail.

Cause

The users are not configured correctly in OpenLDAP.

Corrective action

You can manually fix this issue by using one of the following workarounds.

If your LDAP user object class and group object class are user and group, respectively, perform the following steps:

- a. In the toolbar, click , and then click **Setup > Authentication**.
- b. In the Authentication Service drop-down menu, select Active Directory, and then select Others.
- c. Complete the text fields. If your LDAP user object class and group object class are posixAccount and posixGroup, respectively, perform the following steps:
- d. In the toolbar, click [6], and then click **Setup > Authentication**.
- e. In the Authentication Service drop-down menu, select OpenLDAP, and then select Others.
- f. Complete the text fields. If the first two workarounds do not apply, call the option-set API, and set the auth.ldap.userObjectClass and auth.ldap.groupObjectClass options to the correct values.

Troubleshooting access to CIFS shares

You might not be able to access CIFS shares if the storage objects serving these shares are unavailable. You should review availability events such as Volume Offline, Junction Path Offline, or SVM CIFS Server Down that are generated when these objects are unavailable.

Before you begin

You must have the role of Storage Administrator to perform this task.

About this task

If you have configured an appropriate alert, you will be notified about the availability event through an alert email.

Steps

1. In the **Dashboards/Overview** page, click the appropriate offline event.

For example, if you receive a Volume Offline event, click the **Volume_name Volume Offline** event in the Availability panel in the Unresolved Incidents and Risks area.

2. In the **Event** details page, click **Volume** name in the **Source** field.

- 3. In the **Health/Volume** details page, click the number corresponding to CIFS Shares in the **Related Devices** pane.
- 4. In the Health/Storage Virtual Machine details page, click the CIFS Shares tab.

You can view the number of CIFS shares that are affected.

After you finish

You must resolve the failures by using either OnCommand System Manager or the ONTAP CLI.

Certain special characters do not work with reporting search

Issue

Using the special characters % and _ while searching within a report causes the operation to fail.

Corrective action

If you search for a string that contains % or _, you should use a double backslash before the specified character.

For example, to find a string containing S_10, you should enter S_10.

Monitor and manage cluster performance

Introduction to OnCommand Unified Manager performance monitoring

OnCommand Unified Manager provides performance monitoring capabilities and event root-cause analysis for systems that are running NetApp ONTAP software.

Unified Manager helps you to identify workloads that are overusing cluster components and decreasing the performance of other workloads on the cluster. By defining performance threshold policies you can also specify maximum values for certain performance counters so that events are generated when the threshold is breached. Unified Manager alerts you about these performance events so that you can take corrective action, and bring performance back to normal levels of operation. You can view and analyze events in the Unified Manager UI.

Unified Manager monitors the performance of two types of workloads:

· User-defined workloads

These workloads consist of FlexVol volumes and FlexGroup volumes that you have created in your cluster.

· System-defined workloads

These workloads consist of internal system activity.

Unified Manager performance monitoring features

Unified Manager collects and analyzes performance statistics from systems running ONTAP software. It uses dynamic performance thresholds and user-defined performance thresholds to monitor a variety of performance counters over many cluster components.

A high response time (latency) indicates that the storage object, for example, a volume, is performing slower than normal. This issue also indicates that the performance has decreased for client applications that are using the volume. Unified Manager identifies the storage component where the performance issue lies and provides a list of suggested actions you can take to address the performance issue.

Unified Manager includes the following features:

- Monitors and analyzes workload performance statistics from a system running ONTAP software.
- Tracks performance counters for clusters, nodes, aggregates, ports, SVMs, volumes, LUNs, NVMe namespaces, and LIFs.
- Displays detailed graphs that plot workload activity over time; including IOPS (operations), MBps (throughput), latency (response time), utilization, performance capacity, and cache ratio.
- Enables you to create user-defined performance threshold policies that trigger events and send email alerts when the thresholds are breached.
- Uses system-defined thresholds and dynamic performance thresholds that learn about your workload activity to identify and alert you to performance issues.
- Clearly identifies the cluster component that is in contention.

• Identifies workloads that are overusing cluster components and the workloads whose performance is impacted by the increased activity.

Unified Manager interfaces used to manage storage system performance

There are two user interfaces that OnCommand Unified Manager provides for monitoring and troubleshooting data storage performance issues: the web user interface and the maintenance console.

Unified Manager web UI

The Unified Manager web UI enables an administrator to monitor and troubleshoot storage system issues relating to performance.

This section describes some common workflows that an administrator can follow to troubleshoot storage performance issues displayed in the Unified Manager web UI.

Maintenance console

The maintenance console enables an administrator to monitor, diagnose, and address operating system issues, version upgrade issues, user access issues, and network issues related to the Unified Manager server itself. If the Unified Manager web UI is unavailable, the maintenance console is the only form of access to Unified Manager.

This section provides directions for accessing the maintenance console and using it to resolve issues related to the functioning of the Unified Manager server.

Cluster configuration and performance data collection activity

The collection interval for *cluster configuration data* is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster details in the Unified Manager UI. This interval applies when making changes to a cluster too.

For example, if you add two new volumes to an SVM in a cluster, you see those new objects in the UI after the next polling interval, which could be up to 15 minutes.

Unified Manager collects current *performance statistics* from all monitored clusters every five minutes. It analyzes this data to identify performance events and potential issues. It retains 30 days of five-minute historical performance data and 390 days of one-hour historical performance data. This enables you to view very granular performance details for the current month, and general performance trends for up to a year.

The collection polls are offset by a few minutes so that data from every cluster is not sent at the same time, which could affect performance.

The following table describes the collection activities that Unified Manager performs:

Activity	Time interval	Description
Performance statistics poll	Every 5 minutes	Collects real-time performance data from each cluster.

Activity	Time interval	Description
Statistical analysis	Every 5 minutes	After every statistics poll, Unified Manager compares the collected data against user-defined, system-defined, and dynamic thresholds. If any performance thresholds have been breached, Unified Manager generates events and sends email to specified users, if configured to do so.
Configuration poll	Every 15 minutes	Collects detailed inventory information from each cluster to identify all the storage objects (nodes, SVMs, volumes, and so on).
Summarization	Every hour	Summarizes the latest 12 five- minute performance data collections into hourly averages. The hourly average values are used in some of the UI pages, and they are retained for 390 days.
Forecast analysis and data pruning	Every day after midnight	Analyzes cluster data to establish dynamic thresholds for volume latency and IOPS for the next 24 hours. Deletes from the database any five-minute performance data older than 30 days.
Data pruning	Every day after 2 a.m.	Deletes from the database any events and dynamic thresholds older than 390 days.
Data pruning	Every day after 3:30 a.m.	Deletes from the database any one-hour performance data older than 390 days.

What a data continuity collection cycle is

A data continuity collection cycle retrieves performance data outside of the real-time cluster performance collection cycle that runs, by default, every five minutes. Data continuity collections enable Unified Manager to fill in gaps of statistical data that occur when it was unable to collect real-time data.

Data continuity collection is supported only on clusters installed with ONTAP version 8.3.1 or later software.

Unified Manager performs data continuity collection polls of historical performance data when the following events occur:

· A cluster is initially added to Unified Manager.

Unified Manager gathers historical performance data for the previous 15 days. This enables you to view two weeks of historical performance information for a cluster a few hours after it is added.

Additionally, system-defined threshold events are reported for the previous period, if any exist.



15 days of historical volume statistics are not currently collected.

• The current performance data collection cycle does not finish on time.

If the real-time performance poll goes beyond the five-minute collection period, a data continuity collection cycle is initiated to gather that missing information. Without the data continuity collection, the next collection period is skipped.

- Unified Manager has been inaccessible for a period of time and then it comes back online, as in the following situations:
 - It was restarted.
 - It was shut down during a software upgrade or when creating a backup file.
 - A network outage is repaired.
- A cluster has been inaccessible for a period of time and then it comes back online, as in the following situations:
 - A network outage is repaired.
 - A slow wide area network connection delayed the normal collection of performance data.

A data continuity collection cycle can collect a maximum of 24 hours of historical data. If Unified Manager is down for longer than 24 hours, a gap in performance data appears in the UI pages.

A data continuity collection cycle and a real-time data collection cycle cannot run at the same time. The data continuity collection cycle must finish before the real-time performance data collection is initiated. When the data continuity collection is required to collect more than one hour of historical data, then you see a banner message for that cluster at the top of the Performance dashboard.

What the timestamp means in collected data and events

The timestamp that appears in collected health and performance data, or that appears as the detection time for an event, is based on the ONTAP cluster time, adjusted to the time zone set on the web browser.

It is highly recommended that you use a Network Time Protocol (NTP) server to synchronize the time on your Unified Manager servers, ONTAP clusters, and web browsers.



If you see timestamps that look incorrect for a particular cluster, you might want to check that the cluster time has been set correctly.

Navigating performance workflows in the Unified Manager GUI

The Unified Manager interface provides many pages for the collection and display of performance information. You use the left navigation panel to navigate to pages in the GUI, and you use tabs and links on the pages to view and configure information.

You use all of the following pages to monitor and troubleshoot cluster performance information:

- · dashboard pages
- · storage object inventory pages
- storage object landing pages (including the performance explorer)
- · configuration and setup pages
- · events pages



A page in Unified Manager might display a large amount of information. To see all of the available information, always scroll to the bottom of the page.

Logging in to the UI

You can log in to the Unified Manager UI using a supported web browser.

Before you begin

• The web browser must meet minimum requirements.

See the Interoperability Matrix at mysupport.netapp.com/matrix for the complete list of supported browser versions.

You must have the IP address or URL of the Unified Manager server.

About this task

You are automatically logged out of the session after 24 hours of inactivity.

Steps

- 1. Enter the URL in your web browser, where URL is the IP address or fully qualified domain name (FQDN) of the Unified Manager server:
 - o For IPv4: https://URL/
 - o For IPv6: https://[URL]/ If the server uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate for server authentication.
- 2. At the login screen, enter your user name and password.

If login to the Unified Manager user interface is protected using SAML authentication you will enter your credentials in the identity provider (IdP) login page instead of the Unified Manager login page.

The Dashboards/Overview page is displayed.



If the Unified Manager server is not initialized, a new browser window displays the first experience wizard. You must enter an initial email recipient to which email alerts will be sent, the SMTP server that will handle email communications, and whether AutoSupport is enabled to send information about your Unified Manager installation to technical support. The Unified Manager UI appears after you complete this information.

Graphical interface and navigational paths

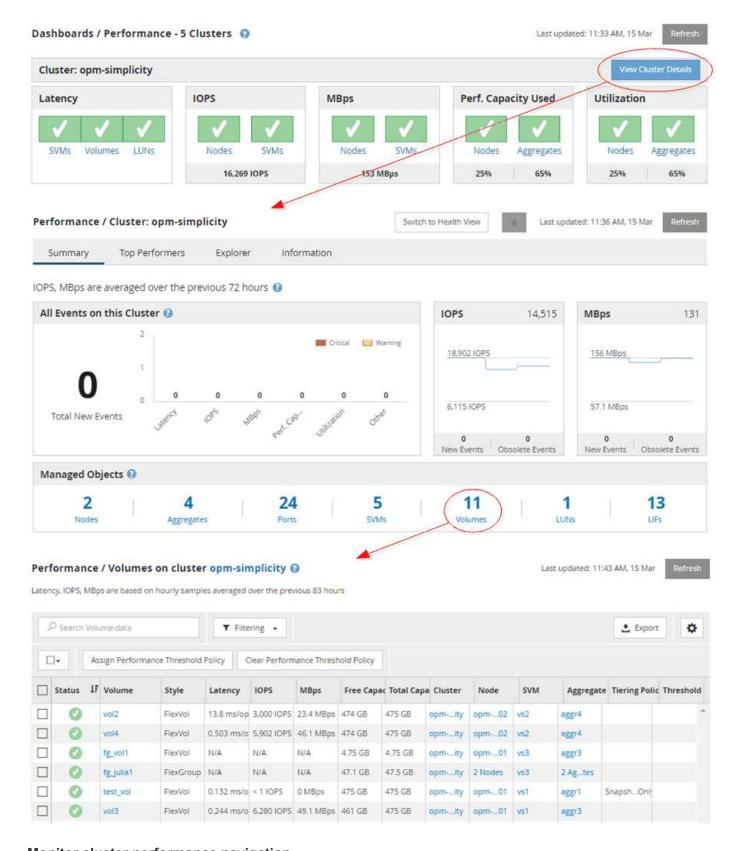
Unified Manager has great flexibility and enables you to accomplish multiple tasks in various ways. There are many navigation paths you will discover as you work in Unified Manager. While not all of the possible combinations of navigations can be shown, you should be familiar with a few of the more common scenarios.

Monitor cluster object navigation

Unified Manager enables you to monitor the performance of all objects in any cluster managed by Unified Manager. Monitoring your storage objects provides you with an overview of cluster and object performance, and includes performance event monitoring. You can view performance and events at a high level, or you can further investigate any details of object performance and performance events.

This is one example of many possible cluster object navigations:

- 1. From the Dashboards/Performance page, identify a cluster you want to investigate and navigate to the selected cluster's landing page.
- 2. From the Performance/Cluster Summary page, identify the cluster object you want to investigate and navigate to that object's inventory page. In this example, **Volumes** is selected to display the Performance/Volumes inventory page.



Monitor cluster performance navigation

Unified Manager enables you to monitor the performance of all clusters managed by Unified Manager. Monitoring your clusters provides you with an overview of cluster and object performance and includes performance event monitoring. You can view performance and events at a high level, or you can further investigate any details of

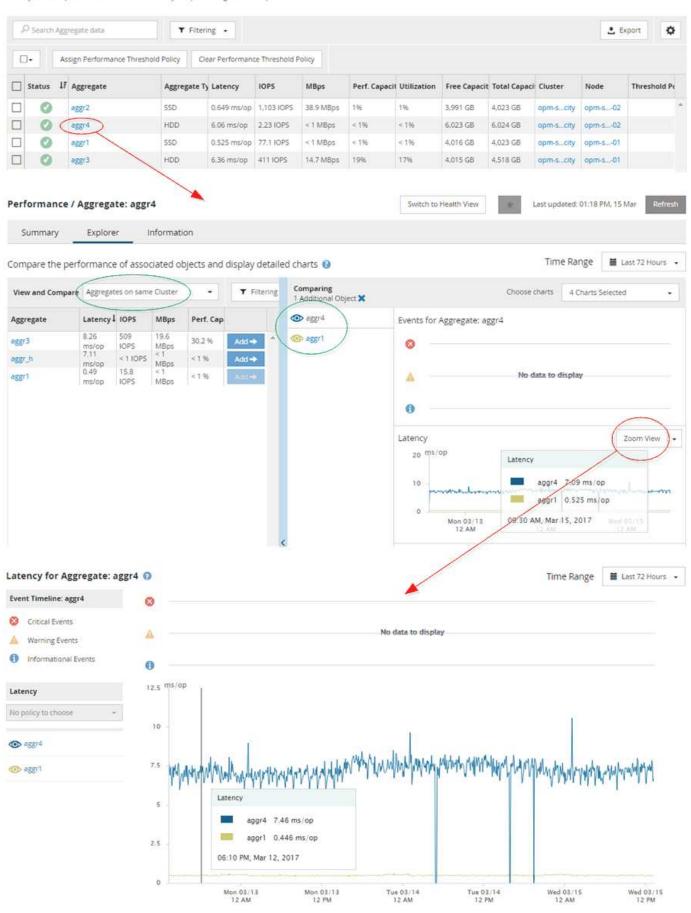
cluster and object performance and performance events.

This is one example of many possible cluster performance navigational paths:

- In the Dashboards/Performance page, identify a cluster you want to investigate and click View Cluster Details to navigate to the selected cluster's landing page.
- 2. From the Performance/Cluster Summary page, identify the object type you want to investigate and click it to view the object inventory page.
 - In this example, **Aggregates** is selected, displaying the Performance/Aggregates inventory page.
- 3. In the Performance/Aggregates page, identify the aggregate you want to investigate and click that aggregate name to navigate to the Performance/Aggregate Explorer page.
- 4. Optionally, select other objects to compare with this aggregate in the View and Compare menu, and then add one of the objects to the comparing pane.
 - Statistics for both objects will appear in the counter charts for comparison.
- 5. In the Comparing pane at the right on the Explorer page, click **Zoom View** in one of the counter charts to view details about the performance history for that aggregate.

Refresh

Latency, IOPS, MBps, Utilization are based on hourly samples averaged over the previous 72 hours



Event investigation navigation

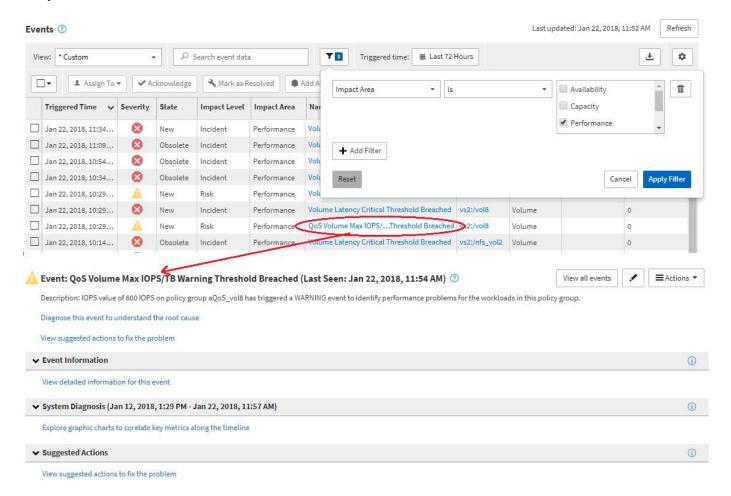
The Unified Manager event detail pages provide you with an in-depth look at any performance event. This is beneficial when investigating performance events, when troubleshooting, and when fine-tuning system performance.

Depending on the type of performance event, you might see one of two types of event detail pages:

- · Event details page for user-defined and system-defined threshold policy events
- · Event details page for dynamic threshold policy events

This is one example of an event investigation navigation.

- 1. In the left navigation pane, click **Events**.
- In the Events inventory page, click the filter button and select **Performance** in the Impact Area to filter the list of events.
- 3. Click the name of the event that you want to investigate and the Event details page is displayed.
- 4. Expand any of the areas, such as Suggested Actions, to view more details about the event that may help you resolve the issue.

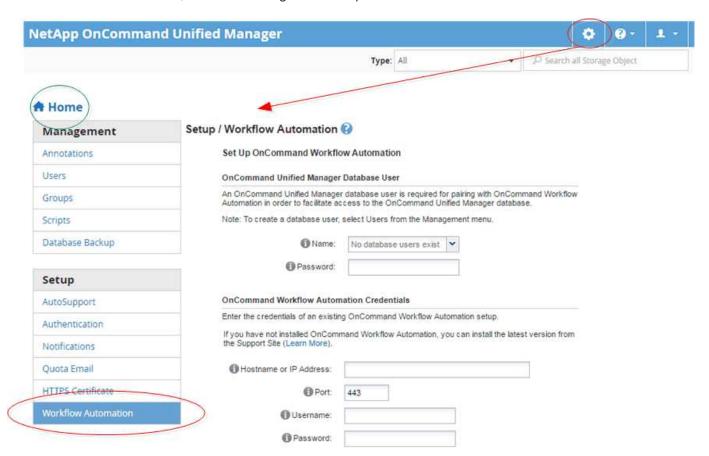


Unified Manager administration navigation

Unified Manager administration functionality enables you to manage users and data sources. You can also accomplish setup tasks such as authentication, AutoSupport,

email, HTTPS certificates, networks, and NTP servers using the Unified Manager Administration page.

This is one example of many possible administration navigational paths. To add or remove a connection to a Workflow Automation server, follow this navigation example:





Click the **Home** icon to return to the main Unified Manager navigation page.

Searching for storage objects

To quickly access a specific object, you can use the **Search all Storage Objects** field at the top-right of the interface. This method of global search across all objects enables you to quickly locate specific objects by type. Search results are sorted by storage object type and you can filter them using the **Type** drop-down menu. A valid search must contain at least three characters.

The global search displays the total number of results, but only the top 20 search results are accessible. Because of this, the global search functionality can be thought of as a shortcut tool for finding specific items if you know the items you want to quickly locate. For complete search results, you can use the search in the object inventory pages and its associated filtering functionality.

You can click the **Type** drop-down box and select **All** to simultaneously search across all objects and events. Alternatively, you can click the **Type** drop-down box to specify the object type. Type any number of characters of the object or event name into the **Search all Storage Objects** field, and then press **Enter** or click **Search All** to display the search results, such as:

Events: performance event IDs

· Clusters: cluster names

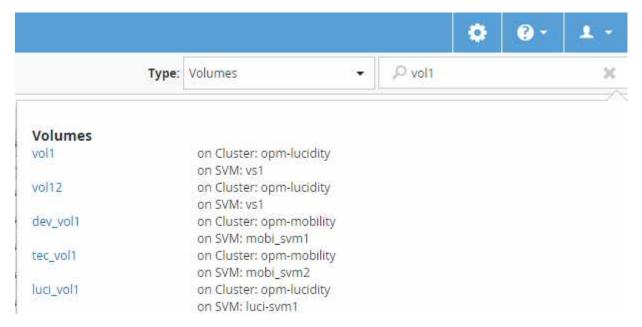
· Nodes: node names

· Aggregates: aggregate names

· SVMs: SVM names

· Volumes: volume names

· LUNs: LUN paths





LIFs and ports are not searchable in the global search bar.

In this example, the **Type** drop-down box has the Volume object type selected. Typing "vol" into the **Search all Storage Objects** field displays a list of all volumes whose names contain these characters. For object searches, you can click any search result to navigate to that object's Performance Explorer page. For event searches, clicking an item in the search result navigates to the Event Details page.



If the search results display several volumes with the same name, the name of the associated clusters and SVMs are not displayed.

Filtering performance inventory page content

You can filter performance inventory data in Unified Manager to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Unified Manager pages to show only the results in which you are interested. This provides a very efficient method of displaying only the performance data in which you are interested.

About this task

Use **Filtering** to customize the grid view based on your preferences. Available filter options are based on the object type being viewed in the grid. If filters are currently applied, an asterisk (*) displays at the left of the Filtering control.

Four types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are contains and starts with .
Number	The operators are greater than and less than .
Resource	The operators are name contains and name starts with.
Status	The operators are is and is not .

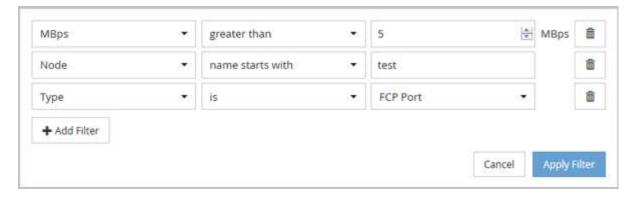
All three fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

You can add filters using the Filtering panel.

- 1. At the top of the page, click **Filtering**. The Filtering panel displays.
- 2. In the Filtering panel, click the left drop-down list, and select an object name: for example, *Cluster*, or a performance counter.
- Click the center drop-down list, and select the boolean operator name contains or name starts with if the
 first selection was an object name. If the first selection was a performance counter, select greater than or
 less than. If the first selection was Status, select is or is not.
- 4. If your search criteria requires a numeric value, up and down arrow buttons display in the field at the right. You can click the up and down arrow buttons to display your desired numeric value.
- 5. If required, type your non-numeric search criteria in the text field at the right.
- To add filters, click Add Filter. An additional filter field displays. Complete this filter using the process
 described in the preceding steps. Note that upon adding your fourth filter, the Add Filter button no longer
 displays.
- 7. Click **Apply Filter**. The filter options are applied to the grid and an asterisk (*) is displayed in the Filtering button.
- 8. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
- 9. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example

The illustration shows the Filtering panel with three filters. The **Add Filter** button displays when you have fewer than the maximum of four filters.



After clicking **Apply Filter**, the Filtering panel closes and applies your filters.

Accessing OnCommand System Manager from the Unified Manager interface

When troubleshooting requires that you make configuration changes to a cluster, you can use the System Manager graphical interface instead of the ONTAP command-line interface. System Manager is included with ONTAP as a web service, it is enabled by default, and it is accessible by using a browser.

Before you begin

You must have a cluster user account configured with the admin role and the http, ontapi, and console application types.

Steps

- 1. In the left navigation pane, click **Dashboards** > **Cluster View**.
- 2. In the Dashboards/Cluster View page, select the cluster that you want to manage.

An overview of the monitoring status, capacity, and performance for that cluster is displayed.

3. Click the System Manager icon.

If the cluster uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate on the cluster for server authentication.

4. Log in to System Manager by using your cluster administrator credentials.

If login to the System Manager user interface is protected using SAML authentication you will enter your credentials in the identity provider (IdP) login page instead of the System Manager login page.

Adding to, and removing storage objects from, the Favorites list

You can add storage objects to a Favorites list so you can monitor the objects for health, capacity, and performance. You can use object status in the Favorites list to determine issues and fix them before they become critical. The Favorites list also provides the most recent monitoring status of a storage object. You can remove storage objects from the Favorites list when you no longer require them to be marked as favorite.

About this task

You can add up to 20 clusters, nodes, aggregates, or volumes to the Favorites list. When you add a node to the Favorites list, it is displayed as a cluster.

Steps

- 1. Go to the **Details** page of the storage object that you want to mark as a favorite.
- 2. Click the star icon () to add the storage object to the Favorites list.

Adding an aggregate to the Favorites list

- 1. In the left navigation pane, click **Health > Aggregates**.
- 2. In the Health/Aggregates inventory page, click the aggregate that you want to add to the Favorites list.
- 3. In the Health/Aggregate details page, click the star icon ().

After you finish

To remove a storage object from the Favorites list, go to the Favorites list page, click the star icon () on the object card you want to remove, and then select the **Remove from Favorites** option.

Bookmarking frequently viewed product pages

You can bookmark frequently accessed product pages from the Unified Manager UI. This enables you to quickly return to these pages. When you view the page later, it displays the latest data.

About this task

You can also copy the link (URL) to the current product page so that you can paste it into an email, or another application, to share it with other people.

Steps

1. Create a bookmark using whatever step is required to bookmark a page in your browser.

The link for the page is saved with details about the page, but you might want to customize the bookmark text to identify the page: for example, "Unified Manager | Node: node-01" or "Unified Manager | User-defined Threshold Event: IOPS volume1".

Bookmarking your favorite Help topics

In the Help Favorites tab, you can bookmark Help topics that you use frequently. Help bookmarks provide fast access to your favorite topics.

Steps

1. Navigate to the Help topic that you want to add as a favorite.

Understanding performance events and alerts

Performance events are notifications that Unified Manager generates automatically when a predefined condition occurs, or when a performance counter value crosses a threshold. Events help you identify performance issues in the clusters that are monitored.

You can configure alerts to send email notification automatically when performance events of certain severity types occur.

Sources of performance events

Performance events are issues related to workload performance on a cluster. They help you identify storage objects with slow response times, also known as high latency. Together with other health events that occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

Unified Manager receives performance events from the following sources:

User-defined performance threshold policy events

Performance issues based on custom threshold values that you have set. You configure performance threshold policies for storage objects; for example, aggregates and volumes, so that events are generated when a threshold value for a performance counter has been breached.

You must define a performance threshold policy and assign it to a storage object to receive these events.

System-defined performance threshold policy events

Performance issues based on threshold values that are system-defined. These threshold policies are included with the installation of Unified Manager to cover common performance problems.

These threshold policies are enabled by default, and you might see events shortly after adding a cluster.

Dynamic performance threshold events

Performance issues that are the result of failures or errors in an IT infrastructure, or from workloads overutilizing cluster resources. The cause of these events might be a simple issue that corrects itself over a period of time or that can be addressed with a repair or configuration change. A dynamic threshold event indicates that volume workloads on an ONTAP system are slow due to other workloads with high usage of shared cluster components.

These thresholds are enabled by default, and you might see events after three days of collecting data from a new cluster.

Performance event severity types

Each performance event is associated with a severity type to help you prioritize the events that require immediate corrective action.

Critical

A performance event occurred that might lead to service disruption if corrective action is not taken immediately.

Critical events are sent from user-defined thresholds only.

Warning

A performance counter for a cluster object is out of normal range and should be monitored to make sure it does not reach the critical severity. Events of this severity do not cause service disruption, and immediate corrective action might not be required.

Warning events are sent from user-defined, system-defined, or dynamic thresholds.

Information

The event occurs when a new object is discovered, or when a user action is performed. For example, when any storage object is deleted or when there are any configuration changes, the event with severity type Information is generated.

Information events are sent directly from ONTAP when it detects a configuration change.

Configuration changes detected by Unified Manager

Unified Manager monitors your clusters for configuration changes to help you determine whether a change might have caused or contributed to a performance event. The Performance Explorer pages display a change event icon () to indicate the date and time when the change was detected.

You can review the performance charts in the Performance Explorer pages and in the Performance/Volume Details page to see whether the change event impacted the performance of the selected cluster object. If the change was detected at or around the same time as a performance event, the change might have contributed to the issue, which caused the event alert to trigger.

Unified Manager can detect the following change events, which are categorized as Informational events:

A volume moves between aggregates.

Unified Manager can detect when the move is in progress, completed, or failed. If Unified Manager is down during a volume move, when it is back up it detects the volume move and displays a change event for it.

• The throughput (MBps or IOPS) limit of a QoS policy group that contains one or more monitored workloads changes.

Changing a policy group limit can cause intermittent spikes in the latency (response time), which might also trigger events for the policy group. The latency gradually returns back to normal and any events caused by the spikes become obsolete.

• A node in an HA pair takes over or gives back the storage of its partner node.

Unified Manager can detect when the takeover, partial takeover, or giveback operation has been completed. If the takeover is caused by a panicked node, Unified Manager does not detect the event.

· An ONTAP upgrade or revert operation is completed successfully.

The previous version and new version are displayed.

What happens when an event is received

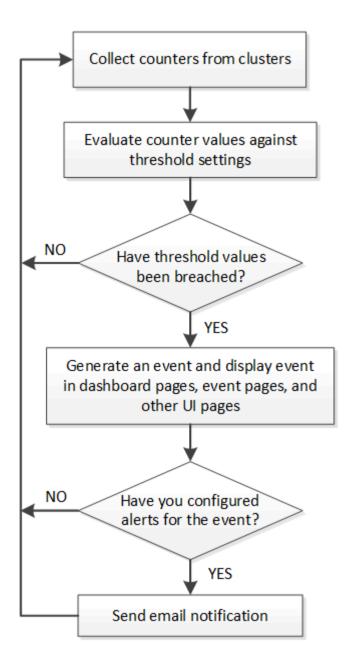
When Unified Manager receives an event, it is displayed in the Dashboards/Overview page, in the Summary and Explorer tabs of the Performance/Cluster page, in the Events inventory page, and in the object-specific inventory page (for example, the Health/Volumes inventory page).

When Unified Manager detects multiple continuous occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events. The duration of the event is incremented to indicate that the event is still active.

Depending on how you configure settings in the Configuration/Alerting page, you can notify other users about these events. The alert causes the following actions to be initiated:

- An email about the event can be sent to all Unified Manager Administrator users.
- The event can be sent to additional email recipients.
- An SNMP trap can be sent to the trap receiver.
- A custom script can be executed to perform an action.

This workflow is shown in the following diagram.



What information is contained in an alert email

Unified Manager alert emails provide the type of event, the severity of the event, the name of the policy that was breached to cause the event, and a description of the event. The email message also provides a hyperlink for each event that enables you to view the details page for the event in the UI.

Alert emails are sent to all users who have subscribed to receive alerts.

If a performance counter or capacity value has a large change during a collection period, it could cause both a critical and a warning event to be triggered at the same time for the same threshold policy. In this case, you may receive one email for the warning event and one for the critical event. This is because Unified Manager enables you to subscribe separately to receive alerts for warning and critical threshold breaches.



After upgrading to Unified Manager 7.2, or greater, links to events and alerts from emails that were send from older versions of Unified Manager will no longer work because of a change in the event and alert URLs.

A sample alert email is shown below:

From: 10.11.12.13@company.com Sent: Tuesday, May 1, 2018 7:45 PM

To: sclaus@company.com; user1@company.com

Subject: Alert from OnCommand Unified Manager: Thin-Provisioned Volume Space At Risk (State: New)

A risk was generated by 10.11.12.13 that requires your attention.

Risk - Thin-Provisioned Volume Space At Risk

Impact Area - Capacity Severity - Warning State - New

Source - svm_n1:/sm_vol_23 Cluster Name - fas3250-39-33-37

Cluster FQDN - fas3250-39-33-37-cm.company.com

Trigger Condition - The thinly provisioned capacity of the volume is 45.73% of the available space on the host aggregate. The capacity of the volume is at risk because of aggregate capacity issues.

Event details:

https://10.11.12.13:443/events/94

Source details:

https://10.11.12.13:443/health/volumes/106

Alert details:

https://10.11.12.13:443/alerting/1

Adding alerts

You can configure alerts to notify you when a particular event is generated. You can configure alerts for a single resource, for a group of resources, or for events of a particular severity type. You can specify the frequency with which you want to be notified and associate a script to the alert.

Before you begin

- You must have configured notification settings such as the user email address, SMTP server, and SNMP trap host to enable the Unified Manager server to use these settings to send notifications to users when an event is generated.
- You must know the resources and events for which you want to trigger the alert, and the user names or email addresses of the users that you want to notify.
- If you want to have a script execute based on the event, you must have added the script to Unified Manager by using the Management/Scripts page.

• You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can create an alert directly from the Event details page after receiving an event in addition to creating an alert from the Configuration/Alerting page, as described here.

Steps

- 1. In the left navigation pane, click **Configuration > Alerting**.
- 2. In the Configuration/Alerting page, click Add.
- 3. In the Add Alert dialog box, click Name, and enter a name and description for the alert.
- 4. Click **Resources**, and select the resources to be included in or excluded from the alert.

You can set a filter by specifying a text string in the **Name contains** field to select a group of resources. Based on the text string that you specify, the list of available resources displays only those resources that match the filter rule. The text string that you specify is case-sensitive.

If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.

- 5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.
 - To select more than one event, press the Ctrl key while you make your selections.
- 6. Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.



If you modify the email address that is specified for the user and reopen the alert for editing, the Name field appears blank because the modified email address is no longer mapped to the user that was previously selected. Also, if you modified the email address of the selected user from the Management/Users page, the modified email address is not updated for the selected user

You can also choose to notify users through SNMP traps.

7. Click Save.

Example of adding an alert

This example shows how to create an alert that meets the following requirements:

- Alert name: HealthTest
- Resources: includes all volumes whose name contains "abc" and excludes all volumes whose name contains "xyz"
- · Events: includes all critical health events
- Actions: includes "sample@domain.com", a "Test" script, and the user has to be notified every 15 minutes

Perform the following steps in the Add Alert dialog box:

- 1. Click Name, and enter HealthTest in the Alert Name field.
- 2. Click Resources, and in the Include tab, select Volumes from the drop-down list.
 - a. Enter abc in the Name contains field to display the volumes whose name contains "abc".
 - b. Select << All Volumes whose name contains 'abc'>> from the Available Resources area, and move it to the Selected Resources area.
 - c. Click **Exclude**, and enter xyz in the **Name contains** field, and then click **Add**.
- Click Events, and select Critical from the Event Severity field.
- 4. Select All Critical Events from the Matching Events area, and move it to the Selected Events area.
- 5. Click Actions, and enter sample@domain.com in the Alert these users field.
- Select Remind every 15 minutes to notify the user every 15 minutes.

You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.

- 7. In the Select Script to Execute menu, select **Test** script.
- 8. Click Save.

Adding alerts for performance events

You can configure alerts for individual performance events just like any other events received by Unified Manager. Additionally, if you want to treat all performance events alike and have email sent to the same person, you can create a single alert to notify you when any critical or warning performance events are triggered.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The example below shows how to create an event for all critical latency, IOPS, and MBps events. You can use this same methodology to select events from all performance counters, and for all warning events.

Steps

- 1. In the left navigation pane, click **Configuration > Alerting**.
- 2. In the Configuration/Alerting page, click Add.
- In the Add Alert dialog box, click Name, and enter a name and description for the alert.
- 4. Do not select any resources on the **Resources** page.

Because no resources are selected, the alert is applied to all clusters, aggregates, volumes, and so on, for which these events are received.

- 5. Click **Events** and perform the following actions:
 - a. In the Event Severity list, select Critical.

- b. In the Event Name Contains field, enter latency and then click the arrow to select all the matching events.
- c. In the Event Name Contains field, enter iops and then click the arrow to select all the matching events.
- d. In the Event Name Contains field, enter mbps and then click the arrow to select all the matching events.
- Click Actions and then select the name of the user who will receive the alert email in the Alert these users field.
- 7. Configure any other options on this page for issuing SNMP taps and executing a script.
- 8. Click Save.

Types of system-defined performance threshold policies

Unified Manager provides some standard threshold policies that monitor cluster performance and generate events automatically. These policies are enabled by default, and they generate warning or information events when the monitored performance thresholds are breached.



System-defined performance threshold policies are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

If you are receiving unnecessary events from any system-defined performance threshold policies, you can disable individual policies from the Configuration/Manage Events page.

Node threshold policies

The system-defined node performance threshold policies are assigned, by default, to every node in the clusters being monitored by Unified Manager:

Node resources over-utilized

Identifies situations in which a single node is operating above the bounds of its operational efficiency, and therefore potentially affecting workload latencies. This is a warning event.

For nodes installed with ONTAP 8.3.x and earlier software, it does this by looking for nodes that are using more than 85% of their CPU and RAM resources (node utilization) for more than 30 minutes.

For nodes installed with ONTAP 9.0 and later software, it does this by looking for nodes that are using more than 100% of their performance capacity for more than 30 minutes.

Node HA pair over-utilized

Identifies situations in which nodes in an HA pair are operating above the bounds of the HA pair operational efficiency. This is an informational event.

For nodes installed with ONTAP 8.3.x and earlier software, it does this by looking at the CPU and RAM usage for the two nodes in the HA pair. If the combined node utilization of the two nodes exceeds 140% for more than one hour, then a controller failover will impact workload latencies.

For nodes installed with ONTAP 9.0 and later software, it does this by looking at the performance capacity used value for the two nodes in the HA pair. If the combined performance capacity used of the two nodes

exceeds 200% for more than one hour, then a controller failover will impact workload latencies.

· Node disk fragmentation

Identifies situations in which a disk or disks in an aggregate are fragmented, slowing key system services and potentially affecting workload latencies on a node.

It does this by looking at certain read and write operation ratios across all aggregates on a node. This policy might also be triggered during SyncMirror resynchronization or when errors are found during disk scrub operations. This is a warning event.



The "Node disk fragmentation" policy analyzes HDD-only aggregates; Flash Pool, SSD, and FabricPool aggregates are not analyzed.

Aggregate threshold policies

The system-defined aggregate performance threshold policy is assigned by default to every aggregate in the clusters being monitored by Unified Manager.

Aggregate disks over-utilized

Identifies situations in which an aggregate is operating above the limits of its operational efficiency, thereby potentially affecting workload latencies. It identifies these situations by looking for aggregates where the disks in the aggregate are more than 95% utilized for more than 30 minutes. This multicondition policy then performs the following analysis to help determine the cause of the issue:

Is a disk in the aggregate currently undergoing background maintenance activity?

Some of the background maintenance activities a disk could be undergoing are disk reconstruction, disk scrub, SyncMirror resynchronization, and reparity.

- Is there a communications bottleneck in the disk shelf Fibre Channel interconnect?
- o Is there too little free space in the aggregate? A warning event is issued for this policy only if one (or more) of the three subordinate policies are also considered breached. A performance event is not triggered if only the disks in the aggregate are more than 95% utilized.



The "Aggregate disks over-utilized" policy analyzes HDD-only aggregates and Flash Pool (hybrid) aggregates; SSD and FabricPool aggregates are not analyzed.

QoS threshold policies

The system-defined QoS performance threshold policies are assigned to any workload that has a configured ONTAP QoS maximum throughput policy (IOPS, IOPS/TB, or MBps). Unified Manager triggers an event when the workload throughput value is 15% less than the configured QoS value.

QoS Max IOPS or MBps threshold

Identifies volumes and LUNs that have exceeded their QoS maximum IOPS or MBps throughput limit, and that are affecting workload latency. This is a warning event.

When a single workload is assigned to a policy group, it does this by looking for workloads that have exceeded the maximum throughput threshold defined in the assigned QoS policy group during each

collection period for the previous hour.

When multiple workloads share a single QoS policy, it does this by adding the IOPS or MBps of all workloads in the policy and checking that total against the threshold.

QoS Peak IOPS/TB or IOPS/TB with Block Size threshold

Identifies volumes that have exceeded their adaptive QoS peak IOPS/TB throughput limit (or IOPS/TB with Block Size limit), and that are affecting workload latency. This is a warning event.

It does this by converting the peak IOPS/TB threshold defined in the adaptive QoS policy into a QoS maximum IOPS value based on the size of each volume, and then it looks for volumes that have exceeded the QoS max IOPS during each performance collection period for the previous hour.



This policy is applied to volumes only when the cluster is installed with ONTAP 9.3 and later software.

When the "block size" element has been defined in the adaptive QoS policy, the threshold is converted into a QoS maximum MBps value based on the size of each volume. Then it looks for volumes that have exceeded the QoS max MBps during each performance collection period for the previous hour.



This policy is applied to volumes only when the cluster is installed with ONTAP 9.5 and later software.

Managing user-defined performance thresholds

Performance threshold policies enable you to determine the point at which Unified Manager generates an event to inform system administrators about issues that could be impacting workload performance. These threshold policies are known as *user-defined* performance thresholds.

This release supports user-defined, system-defined, and dynamic performance thresholds. With dynamic and system-defined performance thresholds, Unified Manager analyzes the workload activity to determine the appropriate threshold value. With user-defined thresholds, you can define the upper performance limits for many performance counters and for many storage objects.



System-defined performance thresholds and dynamic performance thresholds are set by Unified Manager and are not configurable. If you are receiving unnecessary events from any system-defined performance threshold policies, you can disable individual policies from the Configuration/Manage Events page.

How user-defined performance threshold policies work

You set performance threshold policies on storage objects (for example, on aggregates and volumes) so that an event can be sent to the storage administrator to inform the administrator that the cluster is experiencing a performance issue.

You create a performance threshold policy for a storage object by:

· Selecting a storage object

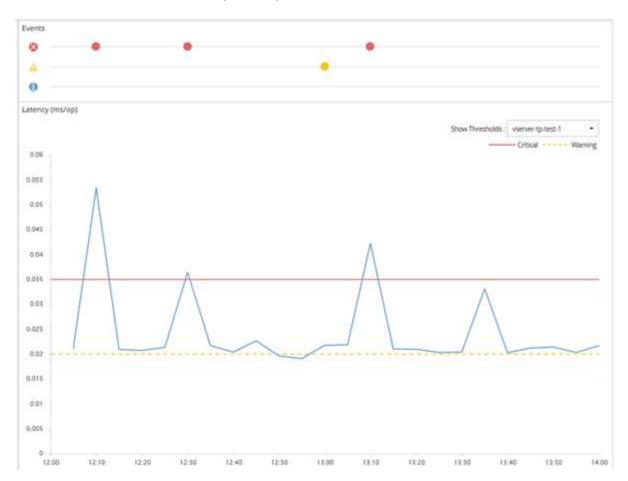
- Selecting a performance counter associated with that object
- Specifying values that define the performance counter upper limits that are considered warning and critical situations
- · Specifying a time period that defines how long the counter must exceed the upper limit

For example, you can set a performance threshold policy on a volume so that you receive a critical event notification whenever IOPS for that volume exceeds 750 operations per second for 10 consecutive minutes. This same threshold policy can also specify that a warning event be sent when IOPS exceeds 500 operations per second for 10 minutes.



The current release provides thresholds that send events when a counter value exceeds the threshold setting. You cannot set thresholds that send events when a counter value falls below a threshold setting.

An example counter chart is shown here, indicating that a warning threshold (yellow icon) was breached at 1:00, and that a critical threshold (red icon) was breached at 12:10, 12:30, and 1:10:



A threshold breach must occur continuously for the specified duration. If the threshold dips below the limit values for any reason, a subsequent breach is considered the start of a new duration.

Some cluster objects and performance counters enable you to create a combination threshold policy that requires two performance counters to exceed their maximum limits before an event is generated. For example, you can create a threshold policy using the following criteria:

Cluster object	Performance counter	Warning threshold	Critical threshold	Duration
Volume	Latency	10 milliseconds	20 milliseconds	15 minutes

Threshold policies that use two cluster objects cause an event to be generated only when both conditions are breached. For example, using the threshold policy defined in the table:

If volume latency is averaging	And aggregate disk utilization is	Then
15 milliseconds	50%	No event is reported.
15 milliseconds	75%	A Warning event is reported.
25 milliseconds	75%	A Warning event is reported.
25 milliseconds	90%	A Critical event is reported.

What happens when a performance threshold policy is breached

When a counter value exceeds its defined performance threshold value for the amount of time specified in the duration, the threshold is breached and an event is reported.

The event causes the following actions to be initiated:

- The event is displayed in the Performance Dashboard, the Performance Cluster Summary page, the Events page, and the object-specific Performance Inventory page.
- (optional) An email alert about the event can be sent to one or more email recipients, and an SNMP trap can be sent to a trap receiver.
- (optional) A script can be executed to automatically modify or update storage objects.

The first action is always executed. You configure whether the optional actions are performed in the Configuration/Alerting page page. You can define unique actions depending on whether a Warning or a Critical threshold policy is breached.

After a performance threshold policy breach has occurred on a storage object, no further events are generated for that policy until the counter value goes below the threshold value, at which point the duration resets for that limit. While the threshold continues to be exceeded, the end time of the event is continually updated to reflect that this event is ongoing.

A threshold event captures, or freezes, the information related to severity and policy definition so that unique threshold information displays with the event, even if the threshold policy is modified in the future.

What performance counters can be tracked using thresholds

Some common performance counters, such as IOPS and MBps, can have thresholds set for all storage objects. There are other counters that can have thresholds set for only certain storage objects.

Available performance counters

Storage object	Performance counter	Description
Cluster	IOPS	Average number of input/output operations the cluster processes per second.
MBps	Average number of megabytes of data transferred to and from this cluster per second.	Node
IOPS	Average number of input/output operations the node processes per second.	MBps
Average number of megabytes of data transferred to and from this node per second.	Latency	Average number of milliseconds the node takes to respond to application requests.
Utilization	Average percentage of the node's CPU and RAM that is being used.	Performance Capacity Used
Average percentage of performance capacity that is being consumed by the node.	Performance Capacity Used - Takeover	Average percentage of performance capacity that is being consumed by the node, plus the performance capacity of its partner node.
Aggregate	IOPS	Average number of input/output operations the aggregate processes per second.
MBps	Average number of megabytes of data transferred to and from this aggregate per second.	Latency
Average number of milliseconds the aggregate takes to respond to application requests.	Utilization	Average percentage of the aggregate's disks that are being used.
Performance Capacity Used	Average percentage of performance capacity that is being consumed by the aggregate.	Storage Virtual Machine (SVM)
IOPS	Average number of input/output operations the SVM processes per second.	MBps

Storage object	Performance counter	Description
Average number of megabytes of data transferred to and from this SVM per second.	Latency	Average number of milliseconds the SVM takes to respond to application requests.
Volume	IOPS	Average number of input/output operations the volume processes per second.
MBps	Average number of megabytes of data transferred to and from this volume per second.	Latency
Average number of milliseconds the volume takes to respond to application requests.	Cache miss ratio	Average percentage of read requests from client applications that are returned from the volume instead of being returned from cache.
LUN	IOPS	Average number of input/output operations the LUN processes per second.
MBps	Average number of megabytes of data transferred to and from this LUN per second.	Latency
Average number of milliseconds the LUN takes to respond to application requests.	Namespace	IOPS
Average number of input/output operations the namespace processes per second.	MBps	Average number of megabytes of data transferred to and from this namespace per second.
Latency	Average number of milliseconds the namespace takes to respond to application requests.	Port
Bandwidth utilization	Average percentage of the port's available bandwidth that is being used.	MBps
Average number of megabytes of data transferred to and from this port per second.	Logical Interface (LIF)	MBps



Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

What objects and counters can be used in combination threshold policies

Only some performance counters can be used together in combination policies. When primary and secondary performance counters are specified, both performance counters must exceed their maximum limits before an event is generated.

Primary storage object and counter	Secondary storage object and counter
Volume Latency	Volume IOPS
Volume MBps	Aggregate Utilization
Aggregate Performance Capacity Used	Node Utilization
Node Performance Capacity Used	Node Performance Capacity Used - Takeover
LUN Latency	LUN IOPS
LUN MBps	Aggregate Utilization
Aggregate Performance Capacity Used	Node Utilization
Node Performance Capacity Used	Node Performance Capacity Used - Takeover



When a volume combination policy is applied to a FlexGroup volume, instead of to a FlexVol volume, only the "Volume IOPS" and "Volume MBps" attributes can be selected as the secondary counter. If the threshold policy contains one of the node or aggregate attributes, then the policy will not be applied to the FlexGroup volume, and you will receive an error message describing this case. This is because FlexGroup volumes can exist on more than one node or aggregate.

Creating user-defined performance threshold policies

You create performance threshold policies for storage objects so that notifications are sent when a performance counter exceeds a specific value. The event notification identifies that the cluster is experiencing a performance issue.

Before you begin

You must have the OnCommand Administrator role.

About this task

You create performance threshold policies by entering the threshold values on the Create Threshold Policy page. You can create new policies by defining all the policy values in this page, or you can make a copy of an

existing policy and change a the values in the copy (called *cloning*).

Valid threshold values are 0.001 through 10,000,000 for numbers, 0.001-100 for percentages, and 0.001-200 for Performance Capacity Used percentages.



The current release provides thresholds that send events when a counter value exceeds the threshold setting. You cannot set thresholds that send events when a counter value falls below a threshold setting.

Steps

1. In the left navigation pane, select **Configuration > Performance Thresholds**.

The Configuration/Performance Thresholds page is displayed.

2. Click the appropriate button depending on whether you want to build a new policy or if you want to clone a similar policy and modify the cloned version.

То	Click
Create a new policy	Create
Clone an existing policy	Select an existing policy and click Clone

The Create Threshold Policy page or Clone Threshold Policy page is displayed.

- Define the threshold policy by specifying the performance counter threshold values you want to set for specific storage objects:
 - a. Select the storage object type and specify a name and description for the policy.
 - b. Select the performance counter to be tracked and specify the limit values that define Warning and Critical events.

You must define at least one Warning or one Critical limit. You do not need to define both types of limits

c. Select a secondary performance counter, if required, and specify the limit values for Warning and Critical events.

Including a secondary counter requires that both counters exceed the limit values before the threshold is breached and an event is reported. Only certain objects and counters can be configured using a combination policy.

d. Select the duration of time for which the limit values must be breached for an event to be sent.

When cloning an existing policy, you must enter a new name for the policy.

4. Click **Save** to save the policy.

You are returned to the Configuration/Performance Thresholds page. A success message at the top of the page confirms that the threshold policy was created and provides a link to the Inventory page for that object type so that you can apply the new policy to storage objects immediately.

After you finish

If you want to apply the new threshold policy to storage objects at this time, you can click the **Go to object_type now** link to go to the Inventory page.

Assigning performance threshold policies to storage objects

You assign a user-defined performance threshold policy to a storage object so that Unified Manager reports an event if the value of the performance counter exceeds the policy setting.

Before you begin

You must have the OnCommand Administrator role.

The performance threshold policy, or policies, that you want to apply to the object must exist.

About this task

You can apply only one performance policy at a time to an object, or to a group of objects.

You can assign a maximum of three threshold policies to each storage object. When assigning policies to multiple objects, if any of the objects already has the maximum number of policies assigned, Unified Manager performs the following actions:

- Applies the policy to all of the selected objects that have not reached their maximum
- Ignores the objects that have reached the maximum number of policies
- · Displays a message that the policy was not assigned to all objects

Additionally, if some objects do not support the counter being tracked in the threshold policy, the policy is not applied to that object. For example, if you create a "Performance Capacity Used" threshold policy, and then you attempt to assign it to a node that does not have ONTAP 9.0 or later software installed, the policy is not applied to that node.

Steps

1. From the Performance inventory page of any storage object, select the object or objects to which you want to assign a threshold policy:

To assign thresholds to	Click
A single object	The check box at the left of that object.
Multiple objects	The check box at the left of each object.
All objects on the page	The drop-down box, and choose Select all objects on this page.
All objects of the same type	The drop-down box, and choose Select all objects.

You can use the sorting and filtering functionality to refine the list of objects on the inventory page to make it easier to apply threshold policies to many objects.

2. Make your selection, and then click Assign Performance Threshold Policy.

The Assign Threshold Policy page is displayed, showing a list of threshold policies that exist for that specific type of storage object.

- 3. Click each policy to display the details of the performance threshold settings to verify that you have selected the correct threshold policy.
- 4. After you have selected the appropriate threshold policy, click **Assign Policy**.

A success message at the top of the page confirms that the threshold policy was assigned to the object or objects, and provides a link to the Alerting page so that you can configure alert settings for this object and policy.

After you finish

If you want to have alerts sent over email, or as an SNMP trap, to notify you that a particular performance event has been generated, you must configure the alert settings in the Configuration/Alerting page.

Viewing performance threshold policies

You can view all of the currently defined performance threshold policies from the Configuration/Performance Thresholds page.

About this task

The list of threshold policies is sorted alphabetically by policy name, and it includes policies for all types of storage objects. You can click a column header to sort the policies by that column. If you are looking for a specific policy, use the filter and search mechanisms to refine the list of threshold policies that appear in the inventory list.

You can hover your cursor over the Policy Name and the Condition name to see the configuration details of the policy. Additionally, you can use the provided buttons to create, clone, edit, and delete user-defined threshold policies.

Steps

1. In the left navigation pane, select **Configuration > Performance Thresholds**.

The Configuration/Performance Thresholds page is displayed.

Editing user-defined performance threshold policies

You can edit the threshold settings for existing performance threshold policies. This can be useful if you find that you are receiving too many or too few alerts for certain threshold conditions.

Before you begin

You must have the OnCommand Administrator role.

About this task

You cannot change the policy name or the type of storage object that is being monitored for existing threshold policies.

Steps

1. In the left navigation pane, select **Configuration > Performance Thresholds**.

The Configuration/Performance Thresholds page displays.

2. Select the threshold policy that you want to change and click Edit.

The Edit Threshold Policy page is displayed.

3. Make your changes to the threshold policy and click **Save**.

You are returned to the Configuration/Performance Thresholds page.

Results

After they are saved, changes are updated immediately on all storage objects that use the policy.

After you finish

Depending on the type of changes that you made to the policy, you may want to review the alert settings configured for the objects that use the policy in the Configuration/Alerting page.

Removing performance threshold policies from storage objects

You can remove a user-defined performance threshold policy from a storage object when you no longer want Unified Manager to monitor the value of the performance counter.

Before you begin

You must have the OnCommand Administrator role.

About this task

You can remove only one policy at a time from a selected object.

You can remove a threshold policy from multiple storage objects by selecting more than one object in the list.

Steps

1. From the **inventory** page of any storage object, select one or more objects that have at least one performance threshold policy applied.

To clear thresholds from	Do this
A single object	Select the check box at the left of that object.
Multiple objects	Select the check box at the left of each object.
All objects on the page	Click and select Select all objects on this page.
All objects of the same type	Click □- and select Select all objects .

2. Click Clear Performance Threshold Policy.

The Clear Threshold Policy page displays, showing a list of threshold policies that are currently assigned to the storage objects.

3. Select the threshold policy you want to remove from the objects and click Clear Policy.

When you select a threshold policy, the details of the policy display so that you can confirm that you have selected the appropriate policy.

What happens when a performance threshold policy is changed

If you adjust the counter value or duration of an existing performance threshold policy, the policy change is applied to all storage objects that use the policy. The new setting takes place immediately, and Unified Manager begins to compare performance counter values to the new threshold settings for all newly collected performance data.

If any active events exist for objects that are using the changed threshold policy, the events are marked as obsolete, and the threshold policy begins monitoring the counter as a newly defined threshold policy.

When viewing the counter on which the threshold has been applied in the Counter Charts Detailed View, the critical and warning threshold lines reflect the current threshold settings. The original threshold settings do not appear on this page even if you view historical data when the old threshold setting was in effect.



Because older threshold settings do not appear in the Counter Charts Detailed View, you might see historical events that appear below the current threshold lines.

What happens to performance threshold policies when an object is moved

Because performance threshold policies are assigned to storage objects, if you move an object, all assigned threshold policies remain attached to the object after the move is completed. For example, if you move a volume or LUN to a different aggregate, the threshold policies are still active for the volume or LUN on the new aggregate.

If a secondary counter condition exists for the threshold policy (a combination policy)--for example, if an additional condition is assigned to an aggregate or a node—the secondary counter condition is applied to the new aggregate or node to which the volume or LUN has been moved.

If any new active events exist for objects that are using the changed threshold policy, the events are marked as obsolete, and the threshold policy begins monitoring the counter as a newly defined threshold policy.

A volume move operation causes ONTAP to send an informational change event. A change event icon appears in the Events timeline on the Performance Explorer page and the Performance/Volume Details page to indicate the time when the move operation was completed.



If you move an object to a different cluster, the user-defined threshold policy is removed from the object. If required, you must assign a threshold policy to the object after the move operation is completed. Dynamic and system-defined threshold policies, however, are applied automatically to an object after it has moved to a new cluster.

Threshold policy functionality during HA takeover and giveback

When a takeover or giveback operation occurs in a high-availability (HA) configuration, objects that are moved from one node to the other node retain their threshold policies in the same manner as in the manual move operations. Because Unified Manager checks for cluster configuration changes every 15 minutes, the impact of the switchover to the new node is not identified until the next poll of the cluster configuration.



If both a takeover and giveback operation occur within the 15-minute configuration change collection period, you might not see the performance statistics move from one node to the other node.

Threshold policy functionality during aggregate relocation

If you move an aggregate from one node to another node using the aggregate relocation start command, both single and combination threshold policies are retained on all objects, and the node portion of the threshold policy is applied to the new node.

Threshold policy functionality during MetroCluster switchover

Objects that move from one cluster to another cluster in a MetroCluster configuration do not retain their userdefined threshold policy settings. If required, you can apply threshold policies on the volumes and LUNs that have moved to the partner cluster. After an object has moved back to its original cluster, the user-defined threshold policy is reapplied automatically.

Volume behavior during switchover and switchback

Monitoring cluster performance from the Performance Dashboard

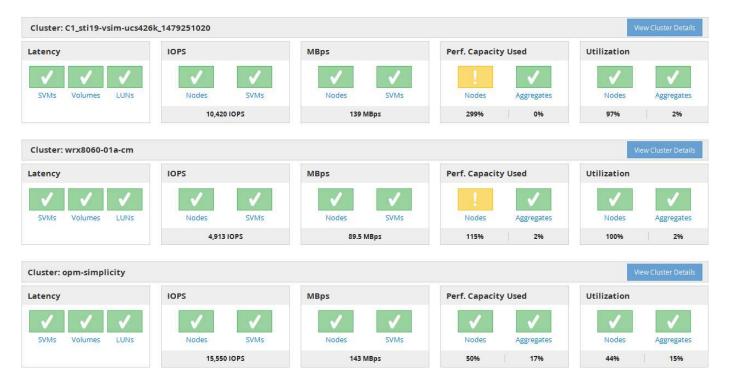
The OnCommand System Manager Performance Dashboard displays the high-level performance status of all clusters being monitored by this instance of Unified Manager. It enables you to assess the overall performance of the managed clusters, and to quickly note, locate, or assign for resolution any specific events identified.

Understanding the Performance dashboard

The Unified Manager Performance dashboard provides a high-level overview of the performance status for all the clusters that are being monitored in your environment.

Clusters that have performance issues are ordered at the top of the page by severity. The information on the dashboard is updated automatically at each five-minute performance collection period.

The following image shows an example of a Unified Manager Performance dashboard that is monitoring two clusters:



The status icons that represent the storage objects can be in the following states, sorted from highest severity to lowest severity:

- Critical (X): One or more new critical performance events have been reported for the object.
- Warning (A): One or more new warning performance events have been reported for the object.
- Normal (
): No new performance events have been reported for the object.



The color indicates whether new events exist for the object. Events that are no longer active, called obsolete events, do not affect the color of the icon.

Cluster performance counters

The following performance categories are displayed for each cluster:

Latency

Shows how quickly the cluster is responding to client application requests, in milliseconds per operation.

• IOPS

Shows the operating speed of the cluster, in number of input/output operations per second.

• MBps

Shows how much data is being transferred to and from the cluster, in megabytes per second.

Performance Capacity Used

Shows whether any nodes or aggregates are overusing their available performance capacity.

Utilization

Shows whether the resources on any nodes or aggregates are being overused.

To analyze the performance of your cluster and storage objects, you can perform one of the following actions:

- You can click **View Cluster Details** to display the Cluster Landing page, where you can view detailed performance and event information for the selected cluster and storage objects.
- You can click one of the red or yellow status icons of an object to display the Inventory page for that object, where you can view details about the storage object.

For example, clicking a volume icon displays the Performance/Volume inventory page with a list of all the volumes in the selected cluster, sorted from worst performance to best performance.

Performance Dashboard cluster banner messages and descriptions

Unified Manager may display cluster banner messages on the Performance Dashboard to alert you to status issues for a particular cluster.

Banner message	Description	Resolution
No performance data is being collected from cluster cluster_name. Restart Unified Manager to correct this issue.	The Unified Manager collection service has stopped and no performance data is being collected from any clusters.	Restart Unified Manager to correct this issue. If this does not correct the issue, contact technical support.
More than x hour(s) of historical data is being collected from cluster cluster_name. Current data collections will start after all historical data is collected.	A data continuity collection cycle is currently running to retrieve performance data outside of the real-time cluster performance collection cycle.	No action is required. Current performance data will be collected after the data continuity collection cycle is completed. A data continuity collection cycle runs when a new cluster is added or when Unified Manager has been unable to collect current performance data for some reason.

Changing the performance statistics collection interval

The default collection interval for performance statistics is 5 minutes. You can change this interval to 10 or 15 minutes if you find that collections from large clusters are not finishing within the default time. This setting affects the collection of statistics from all clusters that this instance of Unified Manager is monitoring.

Before you begin

You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.

About this task

The issue of performance statistics collections not finishing on time is indicated by the banner messages Unable to consistently collect from cluster <cluster_name> or Data collection is taking too long on cluster <cluster name>.

You should change the collection interval only when required because of a statistics collections issue. Do not change this setting for any other reason.



Changing this value from the default setting of 5 minutes can affect the number and frequency of performance events that Unified Manager reports. For example, system-defined performance thresholds trigger events when the policy is exceeded for 30 minutes. When using 5-minute collections, the policy must be exceeded for six consecutive collections. For 15-minute collections the policy must be exceeded for only two collection periods.

A message at the bottom of the Cluster Data Sources page indicates the current statistical data collection interval.

Steps

1. Log in using SSH as the maintenance user to the Unified Manager host.

The Unified Managermaintenance console prompts are displayed.

- 2. Type the number of the menu option labeled **Performance Polling Interval Configuration**, and then press Enter.
- 3. If prompted, enter the maintenance user password again.
- 4. Type the number for the new polling interval that you want to set, and then press Enter.

After you finish

If you changed the Unified Manager collection interval to 10 or 15 minutes, and you have a current connection to an external data provider (such as Graphite), you must change the data provider transmit interval so that it is equal to, or greater, than the Unified Manager collection interval.

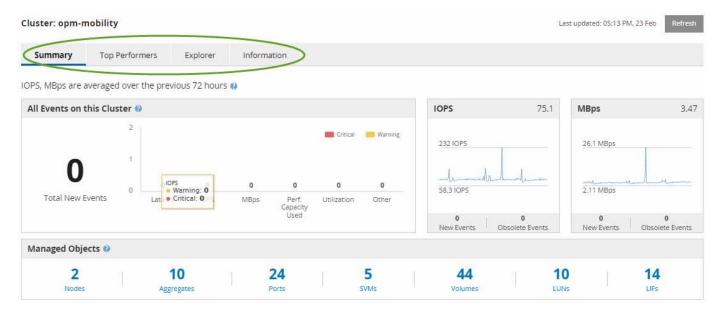
Monitoring cluster performance from the Performance Cluster Landing page

The Performance Cluster Landing page displays the high-level performance status of a selected cluster that is being monitored by an instance of Unified Manager. This page enables you to assess the overall performance of a specific cluster, and to quickly note, locate, or assign for resolution any cluster-specific events that are identified.

Understanding the Performance Cluster Landing page

The Performance Cluster Landing page provides a high-level performance overview of a selected cluster, with an emphasis on the performance status of the top 10 objects within the cluster. Performance issues are displayed at the top of the page, in the All Events on this Cluster panel.

The Performance Cluster Landing page provides a high-level overview of each cluster that is managed by an instance of Unified Manager. This page provides you with information about events and performance, and enables you to monitor and troubleshoot clusters. The following image shows an example of the Performance Cluster Landing page for the cluster called opm-mobility:



The event count on the Cluster Summary page may not match the event count on the Performance Event Inventory page. This is because the Cluster Summary page can show one event each in the Latency and Utilization bars when combination threshold policies have been breached, whereas the Performance Event Inventory page shows only one event when a combination policy has been breached.



If a cluster was removed from being managed by Unified Manager, the status **Removed** is displayed at the right of the cluster name at the top of the page.

Performance Cluster Landing page

The Performance Cluster Landing page displays the high-level performance status of a selected cluster. The page enables you to access complete details of each performance counter for the storage objects on the selected cluster.

You can click the **Favorites** button () to add this object to your list of favorite storage objects. A blue button () indicates that this object is already a favorite.

The Performance Cluster Landing page includes four tabs that separate the cluster details into four areas of information:

- Summary page
 - · Cluster Events pane

- Managed Objects pane
- · Top Performers page
- · Explorer page
- · Information page

Performance Cluster Summary page

The Performance Cluster Summary page provides a summary of the active events, IOPS performance, and MBps performance for a cluster. This page also includes the total count of the storage objects in the cluster.

Cluster performance events pane

The Cluster performance events pane displays performance statistics and all active events for the cluster. This is most helpful when monitoring your clusters and all cluster-related performance and events.

All Events on this Cluster pane

The All Events on this Cluster pane displays all active cluster performance events for the preceding 72 hours. The Total Active Events is displayed at the far left; this number represents the total of all New and Acknowledged events for all storage objects in this cluster. You can click the Total Active Events link to navigate to the Events Inventory page, which is filtered to display these events.

The Total Active Events bar graph for the cluster displays the total number of active critical and warning events:

- Latency (total for nodes, aggregates, SVMs, volumes, LUNs, and namespaces)
- IOPS (total for clusters, nodes, aggregates, SVMs, volumes, LUNs, and namespaces)
- MBps (total for clusters, nodes, aggregates, SVMs, volumes, LUNs, namespaces, ports, and LIFs)
- Performance Capacity Used (total for nodes and aggregates)
- Utilization (total for nodes, aggregates, and ports)
- · Other (cache miss ratio for volumes)

The list contains active performance events triggered from user-defined threshold policies, system-defined threshold policies, and dynamic thresholds.

Graph data (vertical counter bars) is displayed in red () for critical events, and yellow () for warning events. Position your cursor over each vertical counter bar to view the actual type and number of events. You can click **Refresh** to update the counter panel data.

You can show or hide critical and warning events in the Total Active Events performance graph by clicking the **Critical** and **Warning** icons in the legend. If you hide certain event types, the legend icons are displayed in gray.

Counter panels

The counter panels display cluster activity and performance events for the preceding 72 hours, and includes the following counters:

IOPS counter panel

IOPS indicates the operating speed of the cluster in number of input/output operations per second. This counter panel provides a high-level overview of the cluster's IOPS health for the preceding 72-hour period. You can position your cursor over the graph trend line to view the IOPS value for a specific time.

MBps counter panel

MBps indicates how much data has been transferred to and from the cluster in megabytes per second. This counter panel provides a high-level overview of the cluster's MBps health for the preceding 72-hour period. You can position your cursor over the graph trend line to view the MBps value for a specific time.

The number at the top right of the chart in the gray bar is the average value from the last 72-hour period. Numbers shown at the bottom and top of the trend line graph are the minimum and maximum values for the last 72-hour period. The gray bar below the chart contains the count of active (new and acknowledged) events and obsolete events from the last 72-hour period.

The counter panels contain two types of events:

Active

Indicates that the performance event is currently active (new or acknowledged). The issue causing the event has not corrected itself or has not been resolved. The performance counter for the storage object remains above the performance threshold.

Obsolete

Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the storage object is no longer above the performance threshold.

For **Active Events**, if there is one event, you can position your cursor over the event icon and click the event number to link to the appropriate Event Details page. If there is more than one event, you can click **View all Events** to display the Events Inventory page, which is filtered to show all events for the selected object counter type.

Managed Objects pane

The Managed Objects pane in the Performance Summary tab provides a top-level overview of the storage object types and counts for the cluster. This pane enables you to track the status of the objects in each cluster.

The managed objects count is point-in-time data as of the last collection period. New objects are discovered at 15-minute intervals.

Clicking the linked number for any object type displays the object performance inventory page for that object type. The object inventory page is filtered to show only the objects on this cluster.

The managed objects are:

Nodes

A physical system in a cluster.

Aggregates

A set of multiple redundant array of independent disks (RAID) groups that can be managed as a single unit for protection and provisioning.

Ports

A physical connection point on nodes that is used to connect to other devices on a network.

SVMs

A virtual machine providing network access through unique network addresses. An SVM might serve data out of a distinct namespace, and is separately administrable from the rest of the cluster.

Volumes

A logical entity holding accessible user data through one or more of the supported access protocols. The count includes both FlexVol volumes and FlexGroup volumes; it does not include FlexGroup constituents or Infinite Volumes.

• LUNs

The identifier of a Fibre Channel (FC) logical unit or an iSCSI logical unit. A logical unit typically corresponds to a storage volume, and is represented within a computer operating system as a device.

• LIFs

A logical network interface representing a network access point to a node. The count includes all LIF types.

Top Performers page

The Top Performers page displays the storage objects that have the highest performance or the lowest performance, based on the performance counter you select. For example, in the SVMs category, you can display the SVMs that have the highest IOPS, or the highest latency, or the lowest MBps. This page is also shows if any of the top performers have any active performance events (New or Acknowledged).

The Top Performers page displays a maximum of 10 of each object. Note that the Volume object includes both FlexVol volumes and FlexGroup volumes; it does not include FlexGroup constituents or Infinite Volumes.

Time Range

You can select a time range for viewing the top performers; the selected time range applies to all storage objects. Available time ranges:

- Last Hour
- Last 24 Hours
- Last 72 Hours (default)
- Last 7 Days

Metric

Click the **Metric** menu to select a different counter. Counter options are unique to the object type. For example, available counters for the **Volumes** object are **Latency**, **IOPS**, and **MBps**. Changing the counter reloads the panel data with the top performers based on the selected counter.

Available counters:

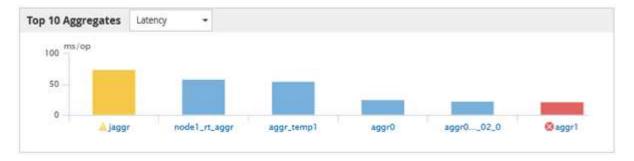
- Latency
- IOPS
- MBps
- Performance Capacity Used (for nodes and aggregates)
- Utilization (for nodes and aggregates)

Sort

Click the **Sort** menu to select an ascending or descending sort for the selected object and counter. The options are **Highest to lowest** and **Lowest to highest**. These options enable you to view the objects with the highest performance or the lowest performance.

Counter bar

The counter bar in the graph shows the performance statistics for each object, represented as a bar for that item. The bar graphs are color-coded. If the counter is not breaching a performance threshold, the counter bar is displayed in blue. If a threshold breach is active (a new or acknowledged event), the bar is displayed in the color for the event: warning events are displayed in yellow (), and critical events are displayed in red (). Threshold breaches are further indicated by severity event indicator icons for warning and critical events.



For each graph, the X axis displays the top performers for the selected object type. The Y axis displays units applicable to the selected counter. Clicking the object name link below each vertical bar graph element navigates to the Performance Landing page for the selected object.

Severity Event indicator

The **Severity Event** indicator icon is displayed at the left of an object name for active critical (**S**) or warning (**A**) events in the top performers graphs. Click the **Severity Event** indicator icon to view:

One event

Navigates to the Event details page for that event.

Two or more events

Navigates to the Event inventory page, which is filtered to display all events for the selected object.

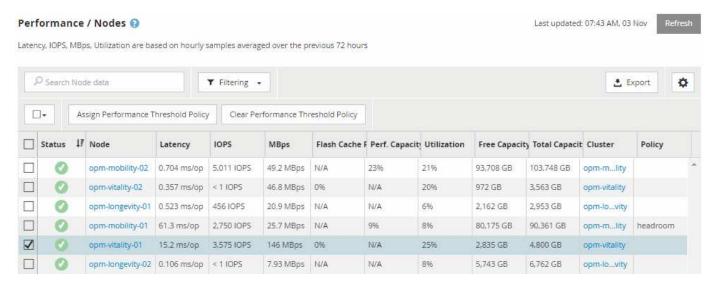
Export button

Creates a .csv file that contains the data that appears in the counter bar. You can choose to create the file for the single cluster you are viewing or for all clusters in the data center.

Monitoring performance using the Performance Inventory pages

The object inventory performance pages display performance information, performance events, and object health for all objects within an object type category. This provides you with an at-a-glance overview of the performance status of each object within a cluster, for example, for all nodes or all volumes.

Object inventory performance pages provide a high-level overview of object status, enabling you to assess the overall performance of all objects and compare object performance data. You can refine the content of object inventory pages by searching, sorting, and filtering. This is beneficial when monitoring and managing object performance, because it enables you to quickly locate objects with performance issues and to begin the troubleshooting process.



By default, objects on the performance inventory pages are sorted based on object performance criticality. Objects with new critical performance events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed. All performance data is based on a 72-hour average.

You can easily navigate from the object inventory performance page to an object details page by clicking the object name in the object name column. For example, on the Performance/Nodes inventory page, you would click a node object in the **Nodes** column. The object details page provides in-depth information and detail about the selected object, including side-by-side comparison of active events.

Object monitoring using the Performance object inventory pages

The Performance object inventory pages enable you to monitor object performance based on the values of specific performance counters or based on performance events. This is beneficial because identifying objects with performance events enables you to investigate the cause of cluster performance issues.

The Performance object inventory pages display the associated counters, associated objects, and performance threshold policies for all objects in all clusters. These pages also enable you to apply performance threshold policies to objects. You can sort the page based on any column, and you can search across all object names or data.

You can export data from these pages to a comma-separated values (.csv) file by using the **Export** button, and then use the exported data to build reports.

Refining Performance inventory page contents

The inventory pages for performance objects contain tools to help you refine object inventory data content, enabling you to locate specific data quickly and easily.

Information contained within the Performance object inventory pages can be extensive, often spanning multiple pages. This kind of comprehensive data is excellent for monitoring, tracking, and improving performance; however, locating specific data requires tools to enable you to quickly locate the data for which you are looking. Therefore, the Performance object inventory pages contain functionality for searching, sorting, and filtering. Additionally, searching and filtering can work together to further narrow your results.

Searching on Object Inventory Performance pages

You can search strings on Object Inventory Performance pages. Use the **Search** field located at the top right of the page to quickly locate data based on either object name or policy name. This enables you to quickly locate specific objects and their associated data, or to quickly locate policies and view associated policy object data.

Steps

1. Perform one of the following options, based on your search requirements:

To locate this	Type this
A specific object	The object name into the Search field, and click Search . The object for which you searched and its related data is displayed.
A user-defined performance threshold policy	All or part of the policy name into the Search field, and click Search . The objects assigned to the policy for which you searched are displayed.

Sorting on the Object Inventory Performance pages

You can sort all data on Object Inventory Performance pages by any column in ascending or descending order. This enables you to quickly locate object inventory data, which is helpful when examining performance or beginning a troubleshooting process.

About this task

The selected column for sorting is indicated by a highlighted column heading name and an arrow icon indicating the sorting direction at the right of the name. An up arrow indicates ascending order; a down arrow indicates descending order. The default sort order is by **Status** (event criticality) in descending order, with the most critical performance events listed first.

Steps

1. You can click a column name to toggle the sort order of the column in ascending or descending order.

The Object Inventory Performance page contents are sorted in ascending or descending order, based on the selected column.

Filtering data in the Object Inventory Performance pages

You can filter data in the Object Inventory Performance pages to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Object Inventory Performance pages to show only the results you have specified. This provides a very efficient method of displaying only the performance data in which you are interested.

About this task

You can use the Filtering panel to customize the grid view based on your preferences. Available filter options are based on the correlated object type being viewed in the grid. If filters are currently applied, an asterisk (*) displays at the left of the Filtering control.

Four types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are contains and starts with .
Number	The operators are greater than and less than .
Resource	The operators are name contains and name starts with.
Status	The operators are is and is not .

All three fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

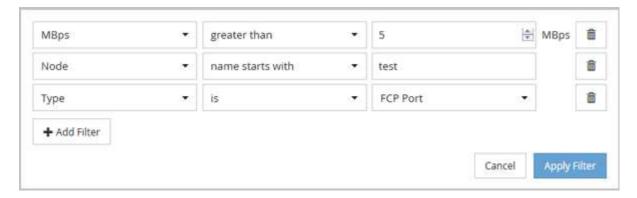
You can add filters using the Filtering panel.

- 1. At the top of the page, click **Filtering**. The Filtering panel displays.
- 2. In the Filtering panel, click the left drop-down list, and select an object name: for example, *Cluster*, or a performance counter.
- 3. Click the center drop-down list, and select the boolean operator **name contains** or **name starts with** if the first selection was an object name. If the first selection was a performance counter, select **greater than** or **less than**. If the first selection was **Status**, select **is** or **is not**.
- 4. If your search criteria requires a numeric value, up and down arrow buttons display in the field at the right. You can click the up and down arrow buttons to display your desired numeric value.
- 5. If required, type your non-numeric search criteria in the text field at the right.

- To add filters, click Add Filter. An additional filter field displays. Complete this filter using the process
 described in the preceding steps. Note that upon adding your fourth filter, the Add Filter button no longer
 displays.
- 7. Click **Apply Filter**. The filter options are applied to the grid and an asterisk (*) is displayed in the Filtering button.
- 8. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
- 9. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example

The illustration shows the Filtering panel with three filters. The **Add Filter** button displays when you have fewer than the maximum of four filters.



After clicking **Apply Filter**, the Filtering panel closes and applies your filters.



Understanding the Unified Manager recommendations to tier data to the cloud

The Performance/Volumes inventory page displays information related to the size of the user data stored on the volume that is inactive (cold). In some cases, Unified Manager identifies certain volumes that would benefit by tiering the inactive data to the cloud tier (cloud provider or StorageGRID) of a FabricPool-enabled aggregate.



FabricPool was introduced in ONTAP 9.2, so if you are using a version of ONTAP software prior to 9.2, the Unified Manager recommendation to tier data requires upgrading your ONTAP software. Additionally, the auto tiering policy was introduced in ONTAP 9.4, so if the recommendation is to use the auto tiering policy, you must upgrade to ONTAP 9.4 or greater.

The following three fields on Performance/Volumes inventory page provide information about whether you can improve your storage system's disk utilization and save space on the performance tier by moving inactive data to the cloud tier.

Tiering Policy

The tiering policy determines whether the data on the volume remains on the performance tier or whether some of the data is moved from the performance tier to the cloud tier.

The value in this field indicates the tiering policy set on the volume, even if the volume does not currently reside on a FabricPool aggregate. The tiering policy takes effect only when the volume is on a FabricPool aggregate.

· Cold Data

The cold data displays the size of the user data stored on the volume that is inactive (cold).

A value is displayed here only when using ONTAP 9.4 or greater software because it requires that the aggregate on which the volume is deployed has the inactive data reporting parameter set to enabled, and that the minimum number of cooling days threshold has been met (for volumes that use the snapshot-only or auto tiering policy). Otherwise the value is listed as "N/A".

Cloud Recommendation

After enough information has been captured about the data activity on the volume, Unified Manager may determine there is no action required, or that you could save space on the performance tier by tiering inactive data to the cloud tier.



The Cold Data field is updated every 15 minutes, but the Cloud Recommendation field is updated every 7 days when the cold data analysis is performed on the volume. Therefore, the exact amount of cold data may differ between the fields. The Cloud Recommendation field displays the date when the analysis was run.

When Inactive Data Reporting is enabled, the Cold Data field displays the exact amount of inactive data. Without the inactive data reporting capability Unified Manager uses performance statistics to determine if data is inactive on a volume. The amount of inactive data is not displayed in the Cold Data field in this case, but it is displayed when you hover your cursor over the word **Tier** to view the cloud recommendation.

The cloud recommendations you will see are:

- Learning. Not enough data has been collected to make a recommendation.
- **Tier**. Analysis has determined that the volume contains inactive (cold) data and that you should configure the volume to move that data to the cloud tier. In some cases this may require that you move the volume to a FabricPool-enabled aggregate first. In other cases where the volume is already on a FabricPool aggregate, you just have to change the tiering policy.
- **No Action**. Either the volume has very little inactive data, the volume is already set to the "auto" tiering policy on a FabricPool aggregate, or the volume is a data protection volume. This value is also displayed when the volume is offline or when it is being used in a MetroCluster configuration.

To move a volume, or to change the volume tiering policy or the aggregate inactive data reporting settings, use OnCommand System Manager, the ONTAP CLI commands, or a combination of these tools.

If you are logged in to Unified Manager with the OnCommand Administrator or Storage Administrator role, the **Configure Volume** link is available in the cloud recommendation when you hover your cursor over the word **Tier**. Click this button to open the Volumes page in System Manager to make the recommended change.

Monitoring performance using the Performance Explorer pages

The Performance Explorer pages display detailed information about the performance of

each object in a cluster. The page provides a detailed view into the performance of all cluster objects, enabling you to select and compare the performance data of specific objects across various time periods.

You can also assess the overall performance of all objects, and compare object performance data in a side-by-side format.

If an object is no longer managed by Unified Manager, the status **Removed** is displayed to the right of the object's name at the top of the Performance Explorer page.

Understanding the root object

The root object is the baseline against which other object comparisons are made. This enables you to view and compare the data from other objects to the root object, providing performance data analysis that helps you to troubleshoot and improve object performance.

The root object name displays at the top of the Comparing pane. Additional objects display below the root object. Although there is no limit to the number of additional objects you can add to the Comparing pane, only one root object is allowed. Data for the root object automatically displays in the graphs in the Counter Charts pane.

You cannot change the root object; it is always set to the object page you are viewing. For example, if you open the Volume Performance Explorer page of Volume1, then Volume1 is the root object and cannot be changed. If you want to compare against a different root object, then you must click the link for an object and open its landing page.



Events and Thresholds are displayed only for root objects.

Apply filtering to reduce the list of correlated objects in the grid

Filtering enables you to display a smaller, more well-defined subset of objects in the grid. For example, if you have 25 volumes in the grid, filtering enables you to view only those volumes that have throughput less than 90 MBps, or latency greater than 1 ms/op.

Specifying a time range for correlated objects

The Time Range selector on the Performance Explorer page enables you to specify the time range for object data comparison. Specifying a time range refines the contents of the Performance Explorer pages to show only the object data within the time range you have specified.

About this task

Refining the time range provides an efficient method of displaying only the performance data in which you are interested. You can select a predefined time range or specify a custom time range. The default time range is the preceding 72 hours.

Selecting a predefined time range

Selecting a predefined time range is a quick and efficient way for you to customize and focus data output when viewing cluster object performance data. When selecting a predefined time range, data for up to 13 months is available.

Steps

- 1. At the top right of the **Performance Explorer** page, click **Time Range**.
- 2. From the right side of the **Time Range Selection** panel, select a predefined time range.
- 3. Click Apply Range.

Specifying a custom time range

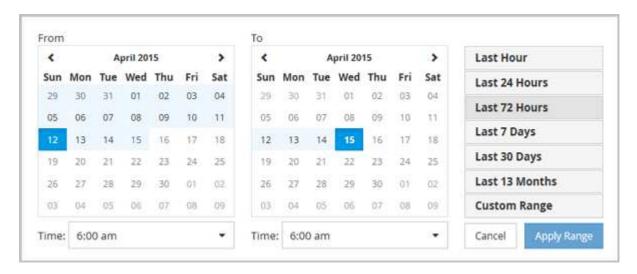
The Performance Explorer page enables you to specify the date and time range for your performance data. Specifying a custom time range provides greater flexibility than using predefined time ranges when refining cluster object data.

About this task

You can select a time range between one hour and 390 days. 13 months equals 390 days because each month is counted as 30 days. Specifying a date and time range provides more detail and enables you to zoom in on specific performance events or series of events. Specifying a time range also aids in troubleshooting potential performance issues, as specifying a date and time range displays data surrounding the performance event in finer detail. Use the **Time Range** control to select predefined date and time ranges, or specify your own custom date and time range of up to 390 days. Buttons for predefined time ranges vary from the **Last Hour** through the **Last 13 Months**.

Selecting the **Last 13 Months** option or specifying a custom date range greater than 30 days displays a dialog box alerting you that performance data displayed for a period greater than 30 days is charted using hourly averages and not 5-minute data polling. Therefore, a loss of timeline visual granularity might occur. If you click the **Do not show again** option in the dialog box, the message does not appear when you select the **Last 13 Months** option or specify a custom date range greater than 30 days. Summary data also applies on a smaller time range, if the time range includes a time/date that is more than 30 days from today.

When selecting a time range (either custom or predefined), time ranges of 30 days or fewer are based on 5-minute interval data samples. Time ranges greater than 30 days are based on one-hour interval data samples.



- 1. Click the **Time Range** drop-down box and the Time Range panel displays.
- 2. To select a predefined time range, click one of the Last... buttons at the right of the Time Range panel. When selecting a predefined time range, data for up to 13 months is available. The predefined time range button you selected is highlighted, and the corresponding days and time display in the calendars and time selectors.
- 3. To select a custom date range, click the start date in the From calendar on the left. Click < or > to navigate forward or backward in the calendar. To specify the end date, click a date in the To calendar on the right. Note that the default end date is today unless you specify a different end date. The Custom Range button at the right of the Time Range panel is highlighted, indicating that you have selected a custom date range.
- 4. To select a custom time range, click the **Time** control below the **From** calendar and select the start time. To specify the end time, click the **Time** control below the **To** calendar on the right and select the end time. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom time range.
- 5. Optionally, you can specify the start and end times when selecting a predefined date range. Select the predefined date range as previously described, then select the start and end times as previously described. The selected dates are highlighted in the calendars, your specified start and end times display in the **Time** controls, and the **Custom Range** button is highlighted.
- 6. After selecting the date and time range, click **Apply Range**. The performance statistics for that time range display in the charts and in the Events timeline .

Defining the list of correlated objects for comparison graphing

You can define a list of correlated objects for data and performance comparison in the Counter Chart pane. For example, if your storage virtual machine (SVM) is experiencing a performance issue, you can compare all volumes in the SVM to identify which volume might be causing the issue.

About this task

You can add any object in the correlated objects grid to the Comparing and Counter Chart panes. This enables you to view and compare data of multiple objects and with the root object. You can add and remove objects to and from the correlated objects grid; however, the root object in the Comparing pane is not removable.



Adding many objects to the Comparing pane may have a negative impact on performance. To maintain performance, you should select a limited number of charts for data comparison.

Steps

1. In the objects grid, locate the object that you want to add, and click the **Add** button.

2. Hide or show data for selected objects:

To do this	Take this action
Hide a selected object	Click the selected object's eye icon () in the Comparing pane. The object's data is hidden, and the eye icon for that object turns gray.
Show a hidden object	Click the gray eye icon of the selected object in the Comparing pane. The eye icon returns to its original color, and the object data is added back into the graphs in the Counter Charts pane.

3. Remove selected objects from the **Comparing** pane:

To do this	Take this action
Remove a selected object	Hover over the selected object's name in the Comparing pane to show the remove object button (X), and then click the button. The object is removed from the Comparing pane, and its data is cleared from the counter charts.
Remove all selected objects	Click the remove all object's button (X) at the top of the Comparing pane. All selected objects and their data are removed, leaving only the root object.

Understanding counter charts

Charts in the Counter Charts pane enable you to view and compare performance data for the root object and for objects you have added from the correlated objects grid. This can help you understand performance trends and isolate and resolve performance issues.

Counter charts displayed by default are Events, Latency,IOPS, and MBps. Optional charts that you can choose to display are Utilization, Performance Capacity Used, Available IOPS, IOPS/TB, and Cache Miss Ratio. Additionally, you can choose to view total values or breakdown values for the Latency, IOPS, MBps, and Performance Capacity Used charts.

The Performance Explorer displays certain counter charts by default; whether the storage object supports them all or not. When a counter is not supported, the counter chart is empty and the message Not applicable for <object> is displayed.

The charts display performance trends for the root object and for all objects you have selected in the Comparing pane. Data in each chart is arranged as follows:

X axis

Displays the specified time period. If you have not specified a time range, the default is the preceding 72-hour period.

Y axis

Displays counter units unique to the selected object, or objects.

Trend line colors match the color of the object name as displayed in the Comparing pane. You can position your cursor over a point on any trend line to view details for time and value for that point.

If you want to investigate a specific period of time within a chart, you can use one of the following methods:

- Use the < button to expand the Counter Charts pane to span the width of the page.
- Use the cursor (when it transitions to a magnifying glass) to select a portion of the timeframe in the chart to focus and enlarge that area. You can click Reset Chart Zoom to return the chart to the default timeframe.
- Use the **Zoom View** button to display a large single counter chart that contains expanded details and threshold indicators.



Occasionally, gaps in the trend lines display. Gaps mean that either Unified Manager failed to collect performance data from the storage system or that Unified Manager might have been down.

Types of performance counter charts

There are standard performance charts that display the counter values for the selected storage object. Each of the Breakdown counter charts display the total values separated out into read, write, and other categories. Furthermore, some Breakdown counter charts display additional detail when the chart is displayed in Zoom view.

The following table shows the available performance counter charts.

Available charts	Chart description
Events	Displays critical, error, warning, and information events in correlation with the statistical charts for the root object. Health events display in addition to performance events to provide a complete picture of the reasons performance may be affected.
Latency - Total	Number of milliseconds required to respond to application requests. Note that the average latency values are I/O weighted.
Latency - Breakdown	The same information shown in Latency Total, but with the performance data separated into read, write, and other latency. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace.
Latency - Cluster Components	The same information shown in Latency Total, but with the performance data separated into latency by cluster component. This chart option applies only when the selected object is a volume.
IOPS - Total	Number of input/output operations processed per second.

Available charts	Chart description
IOPS - Breakdown	The same information shown in IOPS Total, but with the performance data separated into read, write, and other IOPS. When displayed in Zoom view the volumes chart displays QoS minimum and maximum throughput values, if configured in ONTAP. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace.
IOPS - Protocols	The same information shown in IOPS Total, but the performance data is separated into individual charts for CIFS, NFS, FCP, NVMe, and iSCSI protocol traffic. This chart option applies only when the selected object is an SVM.
IOPS/TB - Total	Number of input/output operations processed per second based on the total space that is being consumed by the workload, in terabytes. Also called I/O density, this counter measures how much performance can be delivered by a given amount of storage capacity. When displayed in Zoom view the volumes chart displays QoS expected and peak throughput values, if configured in ONTAP. This chart option applies only when the selected object is a volume.
MBps - Total	Number of megabytes of data transferred to and from the object per second.
MBps - Breakdown	The same information shown in the MBps chart, but with the MBps data separated into disk reads, Flash Cache reads, writes, and other. When displayed in Zoom view, the volumes chart displays QoS maximum throughput values, if configured in ONTAP. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace. Flash Cache data is displayed only for nodes, and only when a Flash Cache module is installed in the node.

Available charts	Chart description
Performance Capacity Used - Total	Percentage of performance capacity that is being consumed by the node or aggregate. Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.
Performance Capacity Used - Breakdown	Performance Capacity Used data separated into user protocols and system background processes. Additionally, the amount of free performance capacity is shown.
Available IOPS - Total	Number of input/output operations per second that are currently available (free) on this object. This number is the result of subtracting the currently used IOPS from the total IOPS that Unified Manager calculates that the object can perform. This chart option applies only when the selected object is a node or aggregate. Available IOPS data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.
Utilization - Total	Available resource percentage of the object that is being used. Utilization indicates node utilization for nodes, disk utilization for aggregates, and bandwidth utilization for ports. This chart option applies only when the selected object is a node, aggregate, or port.
Cache Miss Ratio - Total	Percentage of read requests from client applications that are returned from the disk instead of being returned from the cache. This chart option applies only when the selected object is a volume.

Selecting performance charts to display

The Choose charts drop-down list enables you to select the types of performance counter charts to display in the Counter Charts pane. This enables you to view specific data and counters, based on your performance requirements.

Steps

1. In the Counter Charts pane, click the Choose charts drop-down list.

2. Add or remove charts:

То	Do this
Add or remove individual charts	Click the check boxes next to the charts you want to show or hide
Add all charts	Click Select All
Remove all charts	Click Unselect All

Your chart selections are displayed in the Counter Charts pane. Note that as you add charts, the new charts are inserted into the Counter Charts pane to match the order of the charts listed in the Choose charts drop-down list. Selecting additional charts might require additional scrolling.

Expanding the Counter Charts pane

You can expand the Counter Charts pane so that the charts are larger and more readable.

About this task

After you have defined the comparison objects and the time range for counters, you can view a larger Counter Charts pane. You use the < button in the middle of the Performance Explorer window to expand the pane.

Steps

1. Expand or reduce the **Counter Charts** pane.

То	Do this
Expand the Counter Charts pane to fit the width of the page	Click the < button
Reduce the Counter Charts pane to the right half of the page	Click the > button

Changing the Counter Charts focus to a shorter period of time

You can use your mouse to reduce the time range to focus on a specific period of time in the Counter Chart pane or in the Counter Charts Zoom View window. This enables you to see a more granular and microscopic view of any part of the timeline of performance data, events, and thresholds.

Before you begin

The cursor must have changed to a magnifying glass to indicate that this functionality is active.



When using this feature, which alters the timeline to display values that correspond to the more granular display, the time and date range on the **Time Range** selector does not change from the original values for the chart.

Steps

1. To zoom into a specific period of time, click using the magnifying glass and drag the mouse to highlight the area that you want to see in detail.

The counter values for the time period you select fills the counter chart.

2. To return to the original period of time as set in the **Time Range** selector, click the **Reset Chart Zoom** button.

The counter chart displays in its original state.

Viewing event details in the Events Timeline

You can view all events and their related details in the Events Timeline pane of Performance Explorer. This is a quick and efficient method of viewing all the health and performance events that occurred on the root object during a specified time range, which can be helpful when troubleshooting performance issues.

About this task

The Events Timeline pane shows critical, error, warning, and informational events that occurred on the root object during the selected time range. Each event severity has its own timeline. Single and multiple events are represented by an event dot on the timeline. You can position your cursor over an event dot to see the event details. To increase the visual granularity of multiple events, you can decrease the time range. This spreads out multiple events into single events, enabling you to separately view and investigate each event.

Each performance event dot on the Events Timeline lines up vertically with a corresponding spike in the counter charts trend lines that are displayed below the Events Timeline. This provides a direct visual correlation between events and overall performance. Health events are displayed on the timeline as well, but these types of events do not necessarily line up with a spike in one of the performance charts.

Steps

1. On the **Events Timeline** pane, position the cursor over an event dot on a timeline to view a summary of the event or events at that event point.

A pop-up dialog displays information about the event types, the date and time when the events occurred, the state, and the event duration.

2. View full event details for one event or multiple events:

To do this	Click this
View details for a single event	View Event Detail in the pop-up dialog.

To do this	Click this	
View details for multiple events	View Event Details in the pop-up dialog.	
	Clicking a single event on the multiple events dialog displays the appropriate Event Details page.	

Counter Charts Zoom View

The Counter Charts provide a Zoom View that enables you to zoom in on performance details over your specified time period. This enables you to see performance details and events with much higher granularity, which is beneficial when troubleshooting performance issues.

When displayed in Zoom View, some of the breakdown charts provide additional information than what appears when the chart is not in Zoom View. For example, the IOPS, IOPS/TB, and MBps Breakdown chart Zoom View pages display QoS policy values for volumes and LUNs if they have been set in ONTAP.



For system-defined performance threshold policies, only the "Node resources over-utilized" and "QoS throughput limit breached" policies are available from the **Policies** list. The other system-defined threshold policies are not available at this time.

Displaying the Counter Charts Zoom View

The Counter Charts Zoom View provides a finer level of detail for the selected counter chart and its associated timeline. This magnifies the counter chart data, enabling you to have a sharper view into performance events and their underlying causes.

About this task

You can display the Counter Charts Zoom View for any counter chart.

Steps

- 1. Click **Zoom View** to open the selected chart a new browser window.
- 2. If you are viewing a Breakdown chart and then click **Zoom View** the Breakdown chart is shown in Zoom View. You can select **Total** while in Zoom View if you want to change the view option.

Specifying the time range in Zoom View

The **Time Range** control in the Counter Charts Zoom View window enables you to specify a date and time range for the selected chart. This enables you to quickly locate specific data based on either a preset time range or your own custom time range.

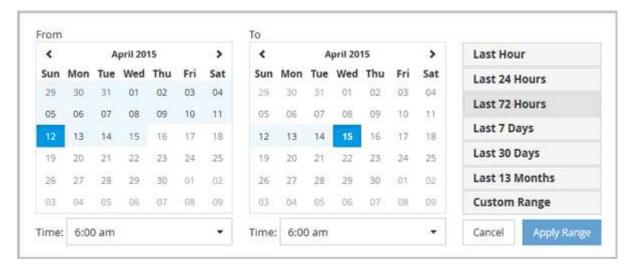
About this task

You can select a time range between one hour and 390 days. 13 months equals 390 days because each month is counted as 30 days. Specifying a date and time range provides more detail and enables you to zoom

in on specific performance events or series of events. Specifying a time range also aids in troubleshooting potential performance issues, as specifying a date and time range displays data surrounding the performance event in finer detail. Use the **Time Range** control to select predefined date and time ranges, or specify your own custom date and time range of up to 390 days. Buttons for predefined time ranges vary from the **Last Hour** through the **Last 13 Months**.

Selecting the **Last 13 Months** option or specifying a custom date range greater than 30 days displays a dialog box alerting you that performance data displayed for a period greater than 30 days is charted using hourly averages and not 5-minute data polling. Therefore, a loss of timeline visual granularity might occur. If you click the **Do not show again** option in the dialog box, the message does not appear when you select the **Last 13 Months** option or specify a custom date range greater than 30 days. Summary data also applies on a smaller time range, if the time range includes a time/date that is more than 30 days from today.

When selecting a time range (either custom or predefined), time ranges of 30 days or fewer are based on 5-minute interval data samples. Time ranges greater than 30 days are based on one-hour interval data samples.



- 1. Click the **Time Range** drop-down box and the Time Range panel displays.
- To select a predefined time range, click one of the Last... buttons at the right of the Time Range panel.
 When selecting a predefined time range, data for up to 13 months is available. The predefined time range
 button you selected is highlighted, and the corresponding days and time display in the calendars and time
 selectors.
- 3. To select a custom date range, click the start date in the From calendar on the left. Click < or > to navigate forward or backward in the calendar. To specify the end date, click a date in the To calendar on the right. Note that the default end date is today unless you specify a different end date. The Custom Range button at the right of the Time Range panel is highlighted, indicating that you have selected a custom date range.
- 4. To select a custom time range, click the **Time** control below the **From** calendar and select the start time. To specify the end time, click the **Time** control below the **To** calendar on the right and select the end time. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom time range.
- 5. Optionally, you can specify the start and end times when selecting a predefined date range. Select the predefined date range as previously described, then select the start and end times as previously described. The selected dates are highlighted in the calendars, your specified start and end times display in the **Time** controls, and the **Custom Range** button is highlighted.
- 6. After selecting the date and time range, click **Apply Range**. The performance statistics for that time range display in the charts and in the Events timeline.

Selecting performance thresholds in Counter Charts Zoom View

Applying thresholds in the Counter Charts Zoom View provides a detailed view of occurrences of performance threshold events. This enables you to apply or remove thresholds, and immediately view the results, which can be helpful while deciding whether troubleshooting should be your next step.

About this task

Selecting thresholds in the Counter Charts Zoom View enables you to view precise data about performance threshold events. You can apply any threshold that appears under the **Policies** area of the Counter Charts Zoom View.

Only one policy at a time can be applied to the object in the Counter Charts Zoom View.

Steps

1. Select or deselect the

that is associated with a policy.

The selected threshold is applied to the Counter Charts Zoom View. Critical thresholds are displayed as a red line; warning thresholds are displayed as a yellow line.

Viewing workload QoS minimum and maximum settings

You can view the ONTAP-defined quality of service (QoS) policy settings on a volume or LUN in the Performance Explorer charts. A throughput maximum setting limits the impact of competing workloads on system resources. A throughput minimum setting ensures that a critical workload meets minimum throughput targets regardless of demand by competing workloads.

About this task

QoS throughput "minimum" and "maximum" IOPS and MBps settings are displayed in the counter charts only if they have been configured in ONTAP. Throughput minimum settings are available only on systems running ONTAP 9.2 or later software, only on AFF systems, and they can be set only for IOPS at this time.

Adaptive QoS policies are available starting with ONTAP 9.3 and are expressed using IOPS/TB instead of IOPS. These policies automatically adjust the QoS policy value based on the volume size, per workload, thereby maintaining the ratio of IOPS to terabytes as the size of the volume changes. You can apply an adaptive QoS policy group to volumes only. The QoS terminology "expected" and "peak" are used for adaptive QoS policies instead of minimum and maximum.

Unified Manager generates warning events for QoS policy breaches when workload throughput has exceeded the defined QoS maximum policy setting during each performance collection period for the previous hour. Workload throughput may exceed the QoS threshold for only a short period of time during each collection period, but Unified Manager displays the "average" throughput during the collection period on the chart. For this reason you may see QoS events while the throughput for a workload might not have crossed the policy threshold shown in the chart.

Steps

1. In the **Performance Explorer** page for your selected volume or LUN, perform the following actions to view

the QoS ceiling and floor settings:

If you want to	Do this
View the IOPS ceiling (the QoS max)	In the IOPS Total or Breakdown chart, click Zoom View .
View the MBps ceiling (the QoS max)	In the MBps Total or Breakdown chart, click Zoom View .
View the IOPS floor (the QoS min)	In the IOPS Total or Breakdown chart, click Zoom View .
View the IOPS/TB ceiling (the QoS peak)	For volumes, in the IOPS/TB chart, click Zoom View .
View the IOPS/TB floor (the QoS expected)	For volumes, in the IOPS/TB chart, click Zoom View .

The dashed, horizontal line indicates the maximum or minimum throughput value set in ONTAP. You can also view when changes to the QoS values were implemented.

2. To view the specific IOPS and MBps values compared to the QoS setting, move your cursor into the chart area to see the popup window.

After you finish

If you notice that certain volumes or LUNs have very high IOPS or MBps and are stressing system resources, you can use System Manager or the ONTAP CLI to adjust the QoS settings so that these workloads do not affect the performance of other workloads.

For more information on adjusting QoS settings, see the ONTAP 9 Performance Monitoring Power Guide.

ONTAP 9 Performance Monitoring Power Guide

How different types of QoS policies are displayed in Unified Manager

You can view the ONTAP-defined quality of service (QoS) policy settings that have been applied to a volume or LUN in the Performance Explorer IOPS, IOPS/TB, and MBps charts. The information displayed in the charts is different depending on the type of QoS policy that has been applied to the workload.

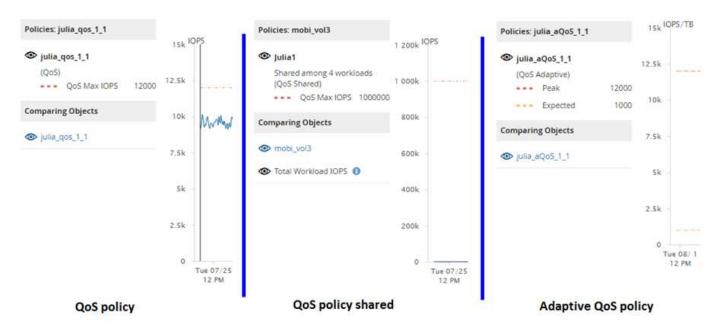
A throughput "ceiling" setting defines the maximum throughput that the workload can consume, and thereby limits the impact on competing workloads for system resources. A throughput "floor" setting defines the minimum throughput that must be available to the workload so that a critical workload meets minimum throughput targets regardless of demand by competing workloads.

Shared and non-shared QoS policies for IOPS and MBps use the terms "minimum" and "maximum" to define the floor and ceiling. Adaptive QoS policies for IOPS/TB, which were introduced in ONTAP 9.3, use the terms "expected" and "peak" to define the floor and ceiling.

While ONTAP enables you to create these two types of QoS policies, depending on how they are applied to workloads there are three ways that the QoS policy will be displayed in the performance charts.

Type of policy	Functionality	Indicator in Unified Manager interface
QoS shared policy assigned to a single workload, or QoS non-shared policy assigned to a single workload or multiple workloads	Each workload can consume the specified throughput setting	Displays "(QoS)"
QoS shared policy assigned to multiple workloads	All workloads share the specified throughput setting	Displays "(QoS Shared)"
Adaptive QoS policy assigned to a single workload or multiple workloads	Each workload can consume the specified throughput setting	Displays "(QoS Adaptive)"

The following figure shows an example of how the three options are shown in the counter charts.



When a normal QoS policy that has been defined in IOPS appears in the IOPS/TB chart for a workload, ONTAP converts the IOPS value to an IOPS/TB value and Unified Manager displays that policy in the IOPS/TB chart along with the text "QoS, defined in IOPS".

When an adaptive QoS policy that has been defined in IOPS/TB appears in the IOPS chart for a workload, ONTAP converts the IOPS/TB value to an IOPS value and Unified Manager displays that policy in the IOPS chart along with the text "QoS Adaptive, defined in IOPS/Used TB" or "QoS Adaptive, defined in IOPS/Allocated TB" depending on how the peak IOPS allocation setting is configured. When the allocation setting is set to "allocated-space", the peak IOPS is calculated based on the size of the volume. When the allocation setting is set to "used-space", the peak IOPS is calculated based on the amount of data stored in the volume, taking into account storage efficiencies.



The IOPS/TB chart displays performance data only when the logical capacity used by the volume is greater than or equal to 1 TB. Gaps are displayed in the chart when the used capacity falls below 1 TB during the selected timeframe.

Viewing volume latency by cluster component

You can view detailed latency information for a volume by using the Performance/Volume Explorer page. The Latency - Total counter chart shows total latency on the volume, and the Latency - Breakdown counter chart is useful for determining the impact of read and write latency on the volume.

About this task

Additionally, the Latency - Cluster Components chart shows a detailed comparison of the latency of each cluster component to help determine how each component contributes to the total latency on the volume. The following cluster components are displayed:

- Network
- QoS Policy
- · Network Processing
- Cluster Interconnect
- · Data Processing
- · Aggregate Operations
- MetroCluster Resources
- Cloud Latency
- Sync SnapMirror

Steps

1. In the **Performance/Volume Explorer** page for your selected volume, from the Latency chart, select **Cluster Components** from the drop-down menu.

The Latency - Cluster Components chart is displayed.

2. To view a larger version of the chart, select **Zoom View**.

The cluster component comparative chart is displayed. You can restrict the comparison by deselecting or selecting the

the thick is associated with each cluster component.

3. To view the specific values, move your cursor into the chart area to see the popup window.

Viewing SVM IOPS traffic by protocol

You can view detailed IOPS information for an SVM by using the Performance/SVM Explorer page. The IOPS - Total counter chart shows total IOPS usage on the SVM, and the IOPS - Breakdown counter chart is useful for determining the impact of read, write, and other IOPS on the SVM.

About this task

Additionally, the IOPS - Protocols chart shows a detailed comparison of the IOPS traffic for each protocol that is being used on the SVM. The following protocols are available:

- CIFS
- NFS
- FCP
- iSCSI
- NVMe

Steps

1. In the **Performance/SVM Explorer** page for your selected SVM, from the IOPS chart, select **Protocols** from the drop-down menu.

The IOPS - Protocols chart is displayed.

2. To view a larger version of the chart, select **Zoom View**.

The IOPS advanced protocol comparative chart is displayed. You can restrict the comparison by deselecting or selecting the that is associated with a protocol.

3. To view the specific values, move your cursor into the chart area of either chart to see the popup window.

Viewing volume and LUN latency charts to verify performance guarantee

You can view the volumes and LUNs that you have subscribed to the "Performance Guarantee" program to verify that latency has not exceeded the level you have been guaranteed.

About this task

The latency performance guarantee is a millisecond per operation value that should not be exceeded. It is based on an hourly average, not on the default five minute performance collection period.

Steps

- 1. In the **Performance Volumes** or **Performance LUNs** inventory page, select the volume or LUN that you are interested in.
- 2. In the **Performance Explorer** page for your selected volume or LUN, choose **Hourly Average** from the **View statistics in** selector.

The horizontal line in the Latency chart will show a smoother line as the five-minute collections are replaced with the hourly average.

3. If you have other volumes on the same aggregate that are under the performance guarantee, you can add those volumes to view their latency value in the same chart.

Components of the Object Landing pages

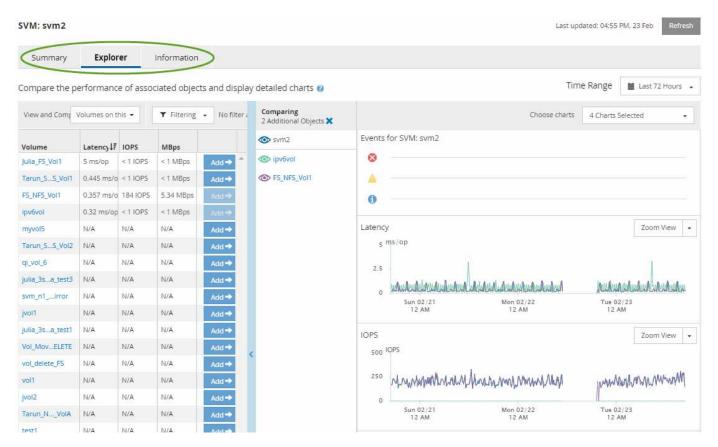
The Object Landing pages provide details about all critical, warning, and informational events. They provide a detailed view into the performance of all cluster objects, enabling you to select and compare individual objects across various time periods.

The Object Landing pages enable you to examine the overall performance of all objects, and to compare object performance data in a side-by-side format. This is beneficial when assessing performance and when troubleshooting events.



The data displayed in the counter summary panels and in the Counter Charts are based on a five-minute sampling interval. The data displayed in the objects inventory grid in the left side of the page is based on a one-hour sampling interval.

The following image shows an example of an Object Landing page displaying the Explorer information:



Depending on the storage object that is being viewed, the Object Landing page can have the following tabs that provide performance data about the object:

Summary

Displays three or four counter charts containing the events and performance per object for the preceding 72-hour period, including a trend line that shows the high and low values during that period.

Explorer

Displays a grid of storage objects that are related to the current object, which enables you to compare the performance values of the current object with those of the related objects. This tab includes up to eleven counter charts and a time range selector, which enable you to perform a variety of comparisons.

Information

Displays values for non-performance configuration attributes about the storage object, including the installed version of ONTAP software, HA partner name, and number of ports and LIFs.

Top Performers

For clusters: Displays the storage objects that have the highest performance or the lowest performance, based on the performance counter that you select.

· Failover Planning

For nodes: Displays the estimate of the performance impact on a node if the HA partner of the node fails.

Details

For volumes: Displays detailed performance statistics for all I/O activity and operations for the selected volume workload. This tab is available for FlexVol volumes, FlexGroup volumes, and constituents of FlexGroups.

Summary page

The Summary page displays counter charts that contain details about the events and performance per object for the preceding 72-hour period. This data is not automatically refreshed, but is current as of the last page load. The charts in the Summary page answer the question *Do I need to look further?*

Charts and counter statistics

The summary charts provide a quick, high-level overview for the last 72-hour period, and help you to identify possible issues that require further investigation.

The Summary page counter statistics are displayed in graphs.

You can position your cursor over the trend line in a graph to view the counter values for a particular point in time. The summary charts also display the total number of active critical and warning events for the preceding 72-hour period for the following counters:

Latency

Average response time for all I/O requests; expressed in milliseconds per operation.

Displayed for all object types.

· IOPS

Average operating speed; expressed in input/output operations per second.

Displayed for all object types.

MBps

Average throughput; expressed in megabytes per second.

Displayed for all object types.

Performance Capacity Used

Percentage of performance capacity that is being consumed by a node or aggregate.

Displayed for nodes and aggregates only. This chart is displayed only when using ONTAP 9.0 or later software.

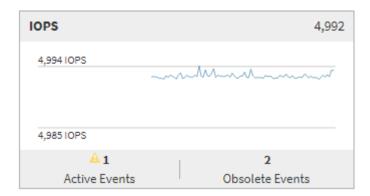
Utilization

Percentage of object utilization for nodes and aggregates, or bandwidth utilization for ports.

Displayed for nodes, aggregates, and ports only.

Positioning the cursor over the event count for Active events shows the type and number of events. Critical events are displayed in red (), and warning events are displayed in yellow ().

The number at the top right of the chart in the gray bar is the average value from the last 72-hour period. Numbers shown at the bottom and top of the trend line graph are the minimum and maximum values for the last 72-hour period. The gray bar below the chart contains the count of active (new and acknowledged) events and obsolete events from the last 72-hour period.



Latency counter chart

The Latency counter chart provides a high-level overview of the object latency for the preceding 72-hour period. Latency refers to the average response time for all I/O requests; expressed in milliseconds per operation, the service time, wait time, or both experienced by a data packet or block in the cluster storage component under consideration.

Top (counter value): The number in the header displays the average for the preceding 72-hour period.

Middle (performance graph): The number at the bottom of the graph displays the lowest latency, and the number at the top of the graph displays the highest latency for the preceding 72-hour period. Position your cursor over the graph trend line to view the latency value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

IOPS counter chart

The IOPS counter chart provides a high-level overview of the object IOPS health for the preceding 72-hour period. IOPS indicates the speed of the storage system in number of input/output operations per second.

Top (counter value): The number in the header displays the average for the preceding 72-hour period.

Middle (performance graph): The number at the bottom of the graph displays the lowest IOPS, and the number at the top of the graph displays the highest IOPS for the preceding 72-hour period. Position your cursor over the graph trend line to view the IOPS value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

MBps counter chart

The MBps counter chart displays the object MBps performance, and indicates how much data has been transferred to and from the object in megabytes per second. The MBps counter chart provides a high-level overview of the object's MBps health for the preceding 72-hour period.

Top (counter value): The number in the header displays the average number of MBps for the preceding 72-hour period.

Middle (performance graph): The value at the bottom of the graph displays the lowest number of MBps, and the value at the top of the graph displays the highest number of MBps for the preceding 72-hour period. Position your cursor over the graph trend line to view the MBps value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

Performance Capacity Used counter chart

The Performance Capacity Used counter chart displays the percentage of performance capacity that is being consumed by the object.

Top (counter value): The number in the header displays the average used performance capacity for the preceding 72-hour period.

Middle (performance graph): The value at the bottom of the graph displays the lowest used performance capacity percentage, and the value at the top of the graph displays the highest used performance capacity percentage for the preceding 72-hour period. Position your cursor over the graph trend line to view the used performance capacity value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

Utilization counter chart

The Utilization counter chart displays the object utilization percentage. The Utilization counter chart provides a high-level overview of the percentage of the object or bandwidth utilization for the preceding 72-hour period.

Top (counter value): The number in the header displays the average utilization percentage for the preceding 72-hour period.

Middle (performance graph): The value at the bottom of the graph displays the lowest utilization percentage, and the value at the top of the graph displays the highest utilization percentage for the preceding 72-hour period. Position your cursor over the graph trend line to view the utilization value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the Active Events link

below the graph to navigate to the Events Inventory page to view complete event details.

Events

The events history table, where applicable, lists the most recent events that occurred on that object. Clicking the event name displays details of the event on the Event Details page.

Components of the Performance Explorer page

The Performance Explorer page enables you to compare the performance of similar objects in a cluster—for example, all the volumes in a cluster. This is beneficial when troubleshooting performance events and fine-tuning object performance. You can also compare objects with the root object, which is the baseline against which other object comparisons are made.

You can click the **Favorites** button () to add this object to your list of favorite storage objects. A blue button () indicates that this object is already a favorite.

You can click the **Switch to Health View** button to display the Health details page for this object. In some cases you can learn important information about the storage configuration settings for this object that may help when troubleshooting an issue.

The Performance Explorer page displays a list of cluster objects and their performance data. This page displays all the cluster objects of the same type (for example, volumes and their object-specific performance statistics) in a tabular format. This view provides an efficient overview of cluster object performance.



If "N/A" appears in any cell of the table, it means that a value for that counter is not available because there is no I/O on that object at this time.

The Performance Explorer page contains the following components:

Time Range

Enables you to select a time range for the object data.

You can choose a predefined range, or specify your own custom time range.

View and Compare

Enables you to select which type of correlated object is displayed in the grid.

The options available depend on the root object type and its available data. You can click the View and Compare drop-down list to select an object type. The object type that you select is displayed in the list.

Filtering

Enables you to narrow the amount of data you receive, based on your preferences.

You can create filters that apply to the object data—for example, IOPS greater than 4. You can add up to four simultaneous filters.

Comparing

Displays a list of the objects that you have selected for comparison with the root object.

Data for the objects in the Comparing pane is displayed in the Counter Charts.

View Statistics In

For volume and LUNs, enables you to select whether the statistics are displayed after each collection cycle (default 5 minutes), or whether the statistics are shown as an hourly average. This functionality enables you to view the latency chart in support of the NetApp "Performance Guarantee" program.

Counter Charts

Displays graphed data for each object performance category.

Typically, only three or four charts are displayed by default. The Choose charts component enables you to display additional charts, or hide specific charts. You can also choose to show or hide the Events Timeline.

Events Timeline

Displays performance and health events occurring across the timeline that you selected in the Time Range component.

Managing performance using performance capacity and available IOPS information

Performance capacity indicates how much throughput you can get out of a resource without surpassing the useful performance of that resource. When viewed using existing performance counters, performance capacity is the point at which you get the maximum utilization from a node or aggregate before latency becomes an issue.

Unified Manager collects performance capacity statistics from nodes and aggregates in each cluster. *Performance capacity used* is the percentage of performance capacity that is currently being used, and *performance capacity free* is the percentage of performance capacity that is still available.

While performance capacity free provides a percentage of the resource that is still available, available IOPS tells you the number of IOPS that can be added to the resource before reaching the maximum performance capacity. By using this metric, you can be sure that you can add workloads of a predetermined number of IOPS to a resource.

Monitoring the performance capacity information has the following benefits:

- Assists with workflow provisioning and balancing.
- Helps you prevent overloading a node or pushing its resources beyond the optimal point, thus reducing the need to troubleshoot.
- Helps you determine with greater precision where additional storage equipment might be needed.

What performance capacity used is

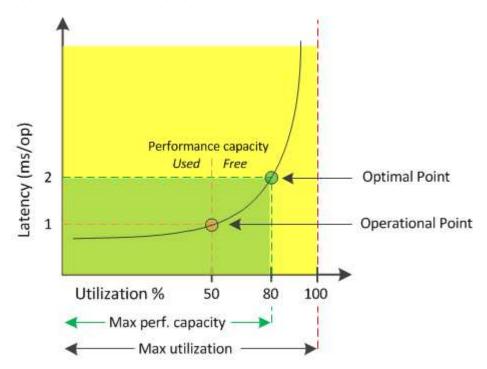
The performance capacity used counter helps you to identify whether the performance of a node or an aggregate is reaching a point where the performance might degrade if the workloads increase. It can also show you if a node or aggregate is currently being

overused during specific periods of time. Performance capacity used is similar to utilization, but the former provides more insight about the available performance capabilities in a physical resource for a specific workload.



Performance capacity data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

The optimal used performance capacity is the point at which a node or an aggregate has optimal utilization and latency (response time), and is being used efficiently. A sample latency versus utilization curve is shown for an aggregate in the following figure.



In this example, the *operational point* identifies that the aggregate is currently operating at 50% utilization with latency of 1.0 ms/op. Based on the statistics captured from the aggregate, Unified Manager determines that additional performance capacity is available for this aggregate. In this example, the *optimal point* is identified as the point when the aggregate is at 80% utilization with latency of 2.0 ms/op. Therefore, you can add more volumes and LUNs to this aggregate so that your systems are used more efficiently.

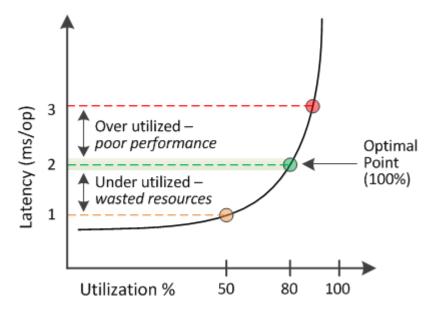
The performance capacity used counter is expected to be a larger number than the "utilization" counter because performance capacity adds in the impact on latency. For example, if a node or aggregate is 70% used, the performance capacity value may be in the 80% to 100% range, depending on the latency value.

In some cases, however, the utilization counter may be higher on the Dashboards/Performance page. This is normal because the dashboard refreshes the current counter values at each collection period; it does not display averages over a period of time like the other pages in the Unified Manager user interface. The performance capacity used counter is best used as an indicator of performance averaged over a period of time, whereas the utilization counter is best used for determining the instantaneous usage of a resource.

What the performance capacity used value means

The performance capacity used value helps you identify the nodes and aggregates that are currently being overutilized or underutilized. This enables you to redistribute workloads in order to make your storage resources more efficient.

The following figure shows the latency versus utilization curve for a resource and identifies, with colored dots, three areas where the current operational point could be located.



• A performance capacity used percentage equal to 100 is at the optimal point.

Resources are being used efficiently at this point.

• A performance capacity used percentage above 100 indicates that the node or aggregate is overutilized, and that workloads are receiving sub-optimal performance.

No new workloads should be added to the resource, and the existing workloads may need to be redistributed.

• A performance capacity used percentage below 100 indicates that the node or aggregate is underutilized, and that resources are not being used effectively.

More workloads can be added to the resource.



Unlike utilization, the performance capacity used percentage can be above 100%. There is no maximum percentage, but resources will typically be in the 110% to 140% range when they are being overutilized. Higher percentages would indicate a resource with serious issues.

What available IOPS is

The available IOPS counter identifies the remaining number of IOPS that can be added to a node or an aggregate before the resource reaches its limit. The total IOPS that a node can provide is based on the physical characteristics of the node—for example, the number of CPUs, the CPU speed, and the amount of RAM. The total IOPS that an aggregate can provide is based on the physical properties of the disks—for example, a SATA, SAS, or SSD disk.

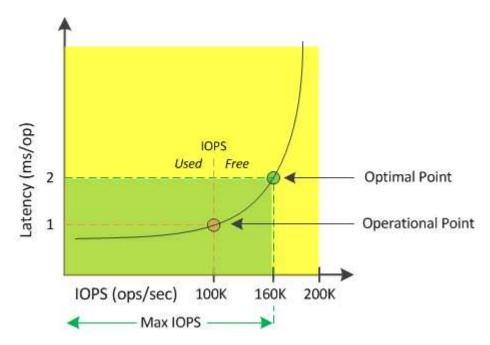
While the performance capacity free counter provides the percentage of a resource that is still available, the available IOPS counter tells you the exact number of IOPS (workloads) can be added to a resource before reaching the maximum performance capacity.

For example, if you are using a pair of FAS2520 and FAS8060 storage systems, a performance capacity free value of 30% means that you have some free performance capacity. However, that value does not provide visibility into how many more workloads you can deploy to those nodes. The available IOPS counter may show that you have 500 available IOPS on the FAS8060, but only 100 available IOPS on the FAS2520.



Available IOPS data is available only when the nodes in a cluster are installed with ONTAP 9.0 or later software.

A sample latency versus IOPS curve for a node is shown in the following figure.



The maximum number of IOPS that a resource can provide is the number of IOPS when the performance capacity used counter is at 100% (the optimal point). The operational point identifies that the node is currently operating at 100K IOPS with latency of 1.0 ms/op. Based on the statistics captured from the node, Unified Manager determines that the maximum IOPS for the node is 160K, which means that there are 60K free or available IOPS. Therefore, you can add more workloads to this node so that your systems are used more efficiently.

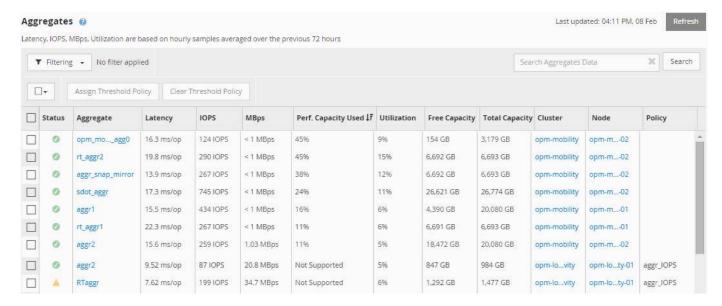


When there is minimal user activity in the resource, the available IOPS value is calculated assuming a generic workload based on approximately 4,500 IOPS per CPU core. This is because Unified Manager lacks the data to accurately estimate the characteristics of the workload being served.

Viewing node and aggregate performance capacity used values

You can monitor the performance capacity used values for all nodes or for all aggregates in a cluster, or you can view details for a single node or aggregate.

Performance capacity used values appear in the Performance Dashboard, Performance Inventory pages, Top Performers page, Create Threshold Policy page, Performance Explorer pages, and in detail charts. For example, the Performance/Aggregate Inventory page provides a column Perf. Capacity Used to view the performance capacity used value for all aggregates.



The status "N/A" is displayed when nodes are not installed with ONTAP 9.0 or later software.

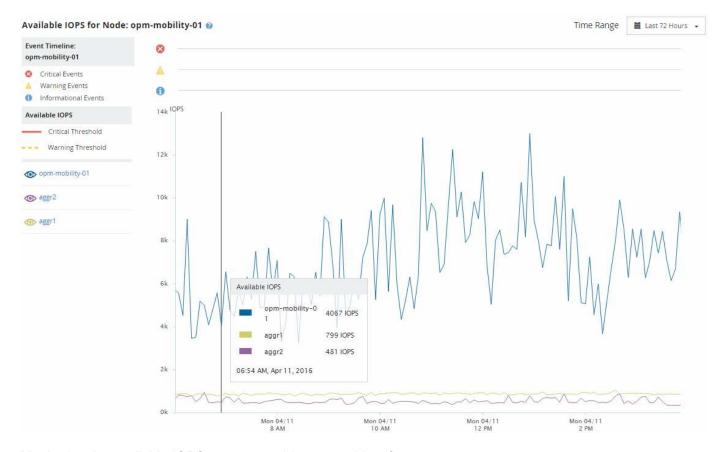
Monitoring the performance capacity used counter enables you to identify the following:

- · Whether any nodes or aggregates on any clusters have a high performance capacity used value
- · Whether any nodes or aggregates on any clusters have active performance capacity used events
- The nodes and aggregates that have the highest and lowest performance capacity used value in a cluster
- Latency and utilization counter values in conjunction with nodes or aggregates that have high performance capacity used values
- · How the performance capacity used values for nodes in an HA pair will be affected if one of the nodes fails
- · The busiest volumes and LUNs on an aggregate that has a high performance capacity used value

Viewing node and aggregate available IOPS values

You can monitor the available IOPS values for all nodes or for all aggregates in a cluster, or you can view details for a single node or aggregate.

Available IOPS values appear in the Performance Explorer page charts. For example, when viewing a node in the Performance/Node Explorer page, you can select the "Available IOPS" counter chart from the list so you can compare the available IOPS values for multiple aggregates on that node.



Monitoring the available IOPS counter enables you to identify:

- The nodes or aggregates that have the greatest available IOPS values to help determine where future workloads can be deployed.
- The nodes or aggregates that have the smallest available IOPS values to identify the resources you should monitor for potential future performance issues.
- The busiest volumes and LUNs on an aggregate that has a small available IOPS value.

Viewing performance capacity counter charts to identify issues

You can view performance capacity used charts for nodes and aggregates on the Performance Explorer page. This enables you to view detailed performance capacity data for the selected nodes and aggregates for a specific timeframe.

About this task

The standard counter chart displays the performance capacity used values for the selected nodes or aggregates. The Breakdown counter chart displays the total performance capacity values for the root object separated into usage based on user protocols versus background system processes. Additionally, the amount of free performance capacity is also shown.

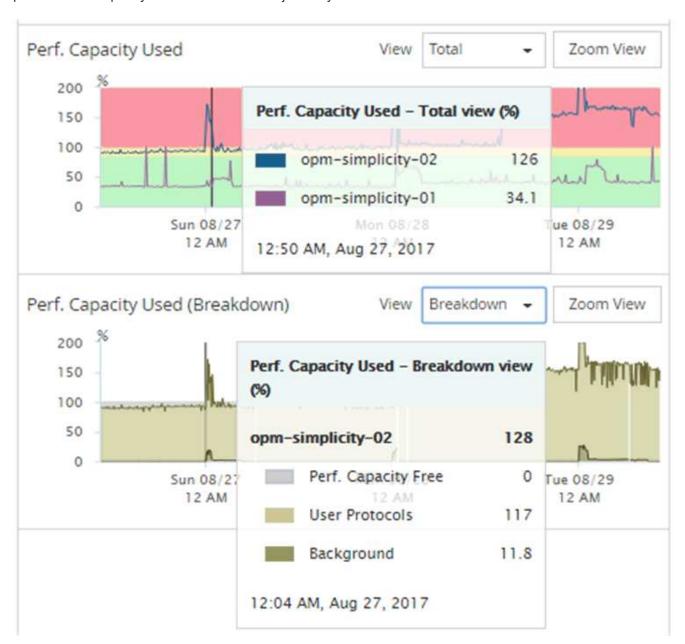


Because some background activities associated with system and data management are identified as user workloads and categorized as user protocols, the user protocols percentage may appear artificially high when those processes run. These processes typically run around midnight when cluster usage is low. If you see a spike in user protocol activity around midnight, verify if cluster backup jobs or other background activities are configured to run at that time.

Steps

- 1. Select the **Explorer** tab from a node or aggregate **Landing** page.
- 2. In the Counter Charts pane, click Choose charts, and then select the Perf. Capacity Used chart.
- Scroll down until you can view the chart.

The colors of the standard chart show when the object is in the optimal range (yellow), when the object is underutilized (green), and when the object is overutilized (red). The Breakdown chart shows detailed performance capacity details for the root object only.



4. If you want to view either chart in a full size format, click **Zoom View**.

In this manner you can open multiple counter charts in a separate windows to compare performance capacity used values with IOPS or MBps values over the same timeframe.

Performance capacity used performance threshold conditions

You can create user-defined performance threshold policies so that events are triggered when the performance capacity used value for a node or aggregate exceeds the defined performance capacity used threshold setting.

Additionally, nodes can be configured with a "Performance capacity used takeover" threshold policy. This threshold policy totals the performance capacity used statistics for both nodes in an HA pair to determine whether either node would lack sufficient capacity if the other node fails. Because the workload during failover is the combination of the two partner nodes' workloads, the same performance capacity used takeover policy can be applied to both nodes.



This performance capacity used equivalency is generally true between nodes. However, if there is significantly more cross-node traffic destined for one of the nodes through its failover partner, the total performance capacity used when running all workloads on one partner node versus the other partner node could be slightly different depending on which node has failed.

The performance capacity used conditions can also be used as secondary performance threshold settings to create a combination threshold policy when defining thresholds for LUNs and volumes. The performance capacity used condition is applied to the aggregate or node on which the volume or LUN resides. For example, you can create a combination threshold policy using the following criteria:

Storage object	Performance counter	Warning threshold	Critical threshold	Duration
Volume	Latency	15 ms/op	25 ms/op	20 minutes

Combination threshold policies cause an event to be generated only when both conditions are breached for the entire duration.

Using the performance capacity used counter to manage performance

Typically, organizations want to operate with a performance capacity used percentage below 100 so that resources are being efficiently used while reserving some additional performance capacity to support peak period demands. You can use threshold policies to customize when alerts are sent for high performance capacity used values.

You can establish specific goals based on your performance requirements. For example, financial services firms might reserve more performance capacity to guarantee the timely execution of trades. These companies might want to set performance capacity used thresholds in the 70-80 percent range. Manufacturing companies with smaller margins might choose to reserve less performance capacity if they are willing to risk performance to better manage IT costs. These companies might set performance capacity used thresholds in the 85-95 percent range.

When the performance capacity used value exceeds the percentage set in a user-defined threshold policy, Unified Manager sends an alert email and adds the event to the Event Inventory page. This enables you to manage potential problems before they impact performance. These events can also be used as indicators that you need to make workload moves and changes within your nodes and aggregates.

Understanding and using the Node Failover Planning page

The Performance/Node Failover Planning page estimates the performance impact on a node if the node's high-availability (HA) partner node fails. Unified Manager bases the estimates on the historical performance of the nodes in the HA pair.

Estimating the performance impact of a failover helps you to plan in the following scenarios:

- If a failover consistently degrades the takeover node's estimated performance to an unacceptable level, you can consider taking corrective actions to reduce the performance impact due to a failover.
- Before initiating a manual failover to perform hardware maintenance tasks, you can estimate how the failover affects the performance of the takeover node in order to determine the best time to perform the task.

Using the Node Failover Planning page to determine corrective actions

Based on the information that is displayed in the Performance/Node Failover Planning page, you can take actions to ensure that a failover does not cause the performance of an HA pair to drop below an acceptable level.

For example, to reduce the estimated performance impact of a failover, you can move some volumes or LUNs from a node in the HA pair to other nodes in the cluster. Doing so ensures that the primary node can continue to deliver acceptable performance after a failover.

Components of the Node Failover Planning page

The components of the Performance/Node Failover Planning page are displayed in a grid and in the Comparing pane. These sections enable you to assess the impact of a node failover on the performance of the takeover node.

Performance statistics grid

The Performance/Node Failover Planning page displays a grid containing statistics for latency, IOPS, utilization, and performance capacity used.



IOPS values displayed in this page and in the Performance/Node Performance Explorer page might not be the same.

In the grid, each node is assigned one of the following roles:

Primary

The node that takes over for the HA partner when the partner fails. The root object is always the Primary node.

Partner

The node that fails in the failover scenario.

· Estimated Takeover

The same as the Primary node. Performance statistics displayed for this node show the takeover node's performance after it takes over the failed partner.



Although the workload of the takeover node is equivalent to the combined workloads of both nodes after a failover, the statistics for the Estimated Takeover node are not the sum of the statistics of the Primary node and the Partner node. For example, if the latency of the Primary node is 2 ms/op and the latency of the Partner node is 3 ms/op, the Estimated Takeover node might have a latency of 4 ms/op. This value is a calculation that Unified Manager performs.

You can click the name of the Partner node if you want it to become the root object. After the Performance/Node Performance Explorer page is displayed, you can click the **Failover Planning** tab to see how performance changes in this node failure scenario. For example, if Node1 is the Primary node and Node2 is the Partner node, you can click Node2 to make it the Primary node. In this way, you can see how the estimated performance changes depending on which node fails.

Comparing pane

The following list describes the components displayed in the Comparing pane by default:

Events charts

They are displayed in the same format as those in the Performance/Node Performance Explorer page. They pertain to the Primary node only.

Counter charts

They display historical statistics for the performance counter shown in the grid. In each chart, the graph for the Estimated Takeover node shows the estimated performance if a failover had occurred at any given time.

For example, suppose the Utilization chart shows 73% for the Estimated Takeover node at 11 a.m. on February 8. If a failover had occurred at that time, the utilization of the takeover node would have been 73%.

The historical statistics help you find the optimal time for initiating a failover, minimizing the possibility of overloading the takeover node. You can schedule a failover only at times when the predicted performance of the takeover node is acceptable.

By default, statistics for both the root object and the partner node are displayed in the Comparing pane. Unlike in the Performance/Node Performance Explorer page, this page does not display the **Add** button for you to add objects for statistics comparison.

You can customize the Comparing pane in the same way as you do in the Performance/Node Performance Explorer page. The following list shows examples of customizing the charts:

- Click a node name to show or hide the node's statistics in the Counter charts.
- Click Zoom View to display a detailed chart for a particular counter in a new window.

Using a threshold policy with the Node Failover Planning page

You can create a node threshold policy so that you can be notified in the Performance/Node Failover Planning page when a potential failover would degrade the

performance of the takeover node to an unacceptable level.

The system-defined performance threshold policy named "Node HA pair over-utilized" generates a warning event if the threshold is breached for six consecutive collection periods (30 minutes). The threshold is considered breached if the combined performance capacity used of the nodes in an HA pair exceeds 200%.

The event from the system-defined threshold policy alerts you to the fact that a failover will cause the latency of the takeover node to increase to an unacceptable level. When you see an event that is generated by this policy for a particular node, you can navigate to the Performance/Node Failover Planning page for that node to view the predicted latency value due to a failover.

In addition to using this system-defined threshold policy, you can create threshold policies by using the "Performance Capacity Used - Takeover" counter, and then apply the policy to selected nodes. Specifying a threshold lower than 200% enables you to receive an event before the threshold for the system-defined policy is breached. You can also specify the minimum period of time for which the threshold is exceeded to less than 30 minutes if you want to be notified before the system-defined policy event is generated.

For example, you can define a threshold policy to generate a warning event if the combined performance capacity used of the nodes in an HA pair exceeds 175% for more than 10 minutes. You can apply this policy to Node1 and Node2, which form an HA pair. After receiving a warning event notification for either Node1 or Node2, you can view the Performance/Node Failover Planning page for that node to assess the estimated performance impact on the takeover node. You can take corrective actions to avoid overloading the takeover node if a failover does happen. If you take action when the combined performance capacity used of the nodes is under 200%, the takeover node's latency does not reach an unacceptable level even if a failover happens during this time.

Using the Performance Capacity Used Breakdown chart for failover planning

The detailed Performance Capacity Used - Breakdown chart shows the performance capacity used for the Primary node and the Partner node. It also shows the amount of free performance capacity on the Estimated Takeover node. This information helps you determine whether you might have a performance issue if the partner node fails.

About this task

In addition to showing the total performance capacity used for the nodes, the Breakdown chart breaks the values for each node into user protocols and background processes.

- User protocols are the I/O operations from user applications to and from the cluster.
- Background processes are the internal system processes involved with storage efficiency, data replication, and system health.

This additional level of detail enables you to determine whether a performance issue is caused by user application activity or background system processes, such as deduplication, RAID reconstruct, disk scrubbing, and SnapMirror copies.

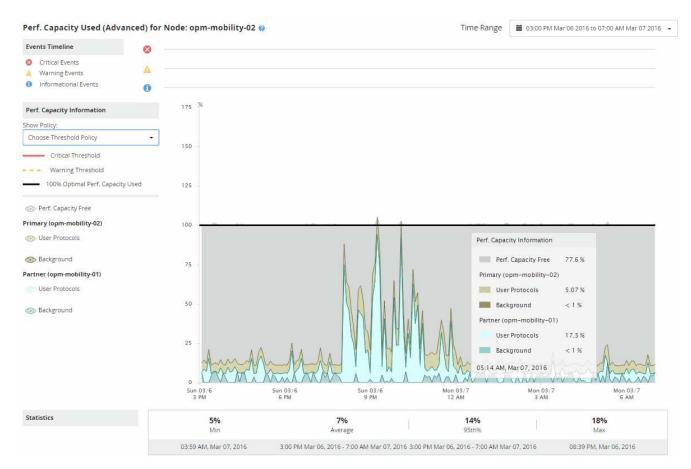
Steps

- 1. Go to the **Performance/Node Failover Planning** page for the node that will serve as the Estimated Takeover node.
- 2. From the **Time Range** selector, choose the period of time for which the historical statistics are displayed in the counter grid and counter charts.

The counter charts with statistics for the Primary node, Partner node, and Estimated Takeover node are displayed.

- 3. From the Choose charts list, select Perf. Capacity Used.
- 4. In the Perf. Capacity Used chart, select Breakdown and click Zoom View.

The detailed chart for Perf. Capacity Used is displayed.



5. Move the cursor over the detailed chart to view the performance capacity used information in the popup window.

The Perf. Capacity Free percentage is the performance capacity available on the Estimated Takeover node. It indicates how much performance capacity is left on the takeover node after a failover. If it is 0%, a failover will cause the latency to increase to an unacceptable level on the takeover node.

Consider taking corrective actions to avoid a low performance capacity free percentage.

If you plan to initiate a failover for node maintenance, choose a time to fail the partner node when the performance capacity free percentage is not at 0.

Collecting data and monitoring workload performance

Unified Manager collects and analyzes workload activity every 5 minutes to identify performance events, and it detects configuration changes every 15 minutes. It retains a maximum of 30 days of 5-minute historical performance and event data, and it uses this data to forecast the expected range for all monitored workloads.



This chapter describes how dynamic thresholds work and how they are used to help monitor workload performance. This chapter is not applicable for statistics or events caused by user-defined or system-defined performance threshold breaches.

Unified Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the expected range for I/O response time and operations can be displayed on the Performance/Volume Details page and in the Event details page. While this activity is being collected, the expected range does not display all changes occurring from workload activity. After collecting 3 days of activity, Unified Manager adjusts the expected range, every 24 hours at 12:00 a.m., to reflect workload activity changes and establish a more accurate performance threshold.

During the first 4 days that Unified Manager is monitoring a volume, if more than 24 hours have passed since the last data collection, the charts on the Performance/Volume Details page will not display the expected range for that volume. Events detected prior to the last collection are still available.



Daylight savings time (DST) changes the system time, which alters the expected range of performance statistics for monitored workloads. Unified Manager immediately begins to correct the expected range, which takes approximately 15 days to complete. During this time you can continue to use Unified Manager, but, since Unified Manager uses the expected range to detect events, some events might not be accurate. Events detected prior to the time change are not affected. Manually changing the time on a cluster, or on a Unified Manager server, to an earlier time will also affect the event analysis results.

Types of workloads monitored by Unified Manager

You can use Unified Manager to monitor the performance of two types of workloads: user-defined and system-defined.

· User-defined workloads

The I/O throughput from applications to the cluster. These are processes involved in read and write requests. A FlexVol volume or FlexGroup volume is a user-defined workload.



Unified Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

If one or more of the following is true for a workload, it cannot be monitored by Unified Manager:

- It is a data protection (DP) copy in read-only mode. (Note that when using ONTAP 8.3 and later, DP volumes are monitored for user-generated traffic.)
- · It is an Infinite Volume.
- It is an offline data clone.
- It is a mirrored volume in a MetroCluster configuration.

· System-defined workloads

The internal processes involved with storage efficiency, data replication, and system health, including:

- Storage efficiency, such as deduplication
- Disk health, which includes RAID reconstruct, disk scrubbing, and so on

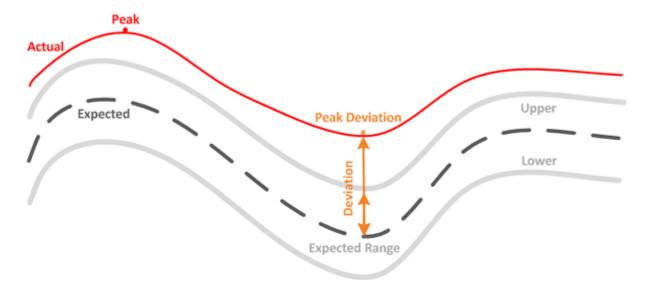
- Data replication, such as SnapMirror copies
- Management activities
- · File system health, which includes various WAFL activities
- File system scanners, such as WAFL scan
- $\,{}^{\circ}\,$ Copy offload, such as offloaded storage efficiency operations from VMware hosts
- System health, such as volume moves, data compression, and so on
- Unmonitored volumes

Performance data for system-defined workloads is displayed in the GUI only when the cluster component used by these workloads is in contention. For example, you cannot search for the name of a system-defined workload to view its performance data in the GUI. If multiple system-defined workloads of the same type are displayed, a letter is appended to the workload name. The letter is intended for use by support personnel.

Workload performance measurement values

Unified Manager measures the performance of workloads on a cluster based on historical and expected statistical values, which form the expected range of values for the workloads. It compares the actual workload statistical values to the expected range to determine when workload performance is too high or too low. A workload that is not performing as expected triggers a performance event report to notify you.

In the following illustration, the actual value, in red, represents the actual performance statistics in the time frame. The actual value has crossed the performance threshold, which is the upper bounds of the expected range. The peak is the highest actual value in the time frame. The deviation measures the change between the expected values and the actual values, while the peak deviation indicates the largest change between the expected values and the actual values.



The following table lists the workload performance measurement values.

Measurement	Description	
Activity	The percentage of the QoS limit used by the workloads in the policy group.	
	If Unified Manager detects a change to a policy group, such as adding or removing a volume or changing the QoS limit, the actual and expected values might exceed 100% of the set limit. If a value exceeds 100% of the set limit it is displayed as >100%. If a value is less than 1% of the set limit it is displayed as <1%.	
Actual	The measured performance value at a specific time for a given workload.	
Deviation	The change between the expected values and the actual values. It is the ratio of the actual value minus the expected value to the upper value of the expected range minus the expected value.	
	A negative deviation value indicates that workload performance is lower than expected, while a positive deviation value indicates that workload performance is higher than expected. If the expected values and the actual value are very low, in the hundredths or thousandths of a percent for example, the deviation will display N/A.	
Expected	The expected values are based on the analysis of historical performance data for a given workload. Unified Manager analyzes these statistical values to determine the expected range of values.	
Expected Range	The expected range of values is a forecast, or prediction, of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Unified Manager triggers a performance event alert.	
Peak	The maximum value measured over a period of time.	
Peak Deviation	The maximum deviation value measured over a period of time.	

Measurement	Description
Queue Depth	The number of pending I/O requests that are waiting at the interconnect component.
Utilization	For the network processing, data processing, and aggregate components, the percentage of busy time to complete workload operations over a period of time. For example, the percentage of time for the network processing or data processing components to process an I/O request or for an aggregate to fulfill a read or write request.
Write Throughput	The amount of write throughput, in Megabytes per second (MBps), from workloads on a local cluster to the partner cluster in a MetroCluster configuration.

What the expected range of performance is

The expected range of values is a forecast, or prediction, of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Unified Manager triggers a performance event alert.

For example, during regular business hours between 9:00 a.m. and 5:00 p.m., most employees might check their email between 9:00 a.m. and 10:30 a.m. The increased demand on the email servers means an increase in workload activity on the back-end storage during this time. Employees might notice slow response time from their email clients.

During the lunch hour between 12:00 p.m. and 1:00 p.m. and at the end of the work day after 5:00 p.m., most employees are likely away from their computers. The demand on the email servers typically decreases, also decreasing the demand on back-end storage. Alternatively, there could be scheduled workload operations, such as storage backups or virus scanning, that start after 5:00 p.m. and increase activity on the back-end storage.

Over several days, the increase and decrease in workload activity determines the expected range of activity, with upper and lower boundaries for a workload. When the actual workload activity for an object is outside the upper or lower boundaries, and remains outside the boundaries for a period of time, this might indicate that the object is being overused or underused.

How the expected range is formed

Unified Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the expected range for I/O response time and operations can be displayed in the GUI. The minimum required data collection does not account for all changes occurring from workload activity. After collecting the first 3 days of activity, Unified Manager adjusts the expected range, every 24 hours at 12:00 a.m., to reflect workload activity changes and establish a more accurate performance threshold.



Daylight savings time (DST) changes the system time, which alters the expected range of performance statistics for monitored workloads. Unified Manager immediately begins to correct the expected range, which takes approximately 15 days to complete. During this time you can continue to use Unified Manager, but, since Unified Manager uses the expected range to detect events, some events might not be accurate. Events detected prior to the time change are not affected. Manually changing the time on a cluster, or on a Unified Manager server, to an earlier time will also affect the event analysis results.

How the expected range is used in performance analysis

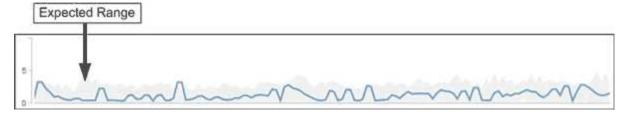
Unified Manager uses the expected range to represent the typical I/O latency (response time) and IOPS (operations) activity for your monitored workloads. It alerts you when the actual latency for a workload is above the upper bounds of the expected range, which triggers a performance event, so that you can analyze the performance issue and take corrective action for resolving it.

The expected range sets the performance baseline for the workload. Over time, Unified Manager learns from past performance measurements to forecast the expected performance and activity levels for the workload. The upper boundary of the expected range establishes the performance threshold. Unified Manager uses the baseline to determine when the actual latency or operations are above or below a threshold, or outside the bounds of their expected range. The comparison between the actual values and the expected values creates a performance profile for the workload.

When the actual latency for a workload exceeds the performance threshold, due to contention on a cluster component, the latency is high and the workload performs more slowly than expected. The performance of other workloads that share the same cluster components might also be slower than expected.

Unified Manager analyzes the threshold crossing event and determines whether the activity is a performance event. If the high workload activity remains consistent for a long period of time, such as several hours, Unified Manager considers the activity to be normal and dynamically adjusts the expected range to form the new performance threshold.

Some workloads might have consistently low activity, where the expected range for the operations or the latency does not have a high rate of change over time. To minimize the number of event alerts, during analysis of performance events, Unified Manager triggers an event only for low-activity volumes whose operations and latencies are much higher than expected.



In this example, the latency for a volume has an expected range, in gray, of 0 milliseconds per operation (ms/op) at its lowest and 5 ms/op at its highest. If the actual latency, in blue, suddenly increases to 10 ms/op, due to an intermittent spike in network traffic or contention on a cluster component, it is then above the expected range and has exceeded the performance threshold.

When network traffic has decreased, or the cluster component is no longer in contention, the latency returns within the expected range. If the latency remains at or above 10 ms/op for a long period of time, you might need to take corrective action to resolve the event.

How Unified Manager uses workload latency to identify performance issues

The workload latency (response time) is the time it takes for a volume on a cluster to respond to I/O requests from client applications. Unified Manager uses the latency to detect and alert you to performance events.

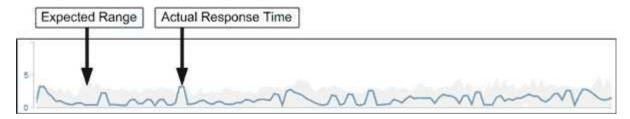
A high latency means that requests from applications to a volume on a cluster are taking longer than usual. The cause of the high latency could be on the cluster itself, due to contention on one or more cluster components. High latency could also be caused by issues outside of the cluster, such as network bottlenecks, issues with the client hosting the applications, or issues with the applications themselves.



Unified Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

Operations on the cluster, such as making backups or running deduplication, that increase their demand of cluster components shared by other workloads can also contribute to high latency. If the actual latency exceeds the performance threshold of the expected range, Unified Manager analyzes the event to determine whether it is a performance event that you might need to resolve. The latency is measured in milliseconds per operation (ms/op).

On the Performance/Volume Details page, you can view an analysis of the latency statistics to see how the activity of individual processes, such as read and write requests, compares to the overall latency statistics. The comparison helps you determine which operations have the highest activity or whether specific operations have abnormal activity that is impacting the latency for a volume. When analyzing performance events, you can use the latency statistics to determine whether an event was caused by an issue on the cluster. You can also identify the specific workload activities or cluster components that are involved in the event.



This example shows the Latency chart on the Performance/Volume Details page. The actual response time (latency) activity is a blue line and the expected range is gray.

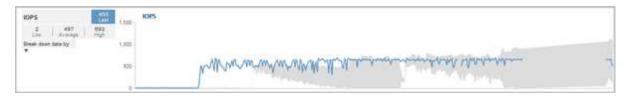


There can be gaps in the blue line if Unified Manager was unable to gather data. This can occur because the cluster or volume was unreachable, Unified Manager was turned off during that time, or the collection was taking longer than the 5 minute collection period.

How cluster operations can affect workload latency

Operations (IOPS) represent the activity of all user-defined and system-defined workloads on a cluster. The IOPS statistics help you determine whether cluster processes, such as making backups or running deduplication, are impacting workload latency (response time) or might have caused, or contributed to, a performance event.

When analyzing performance events, you can use the IOPS statistics to determine whether a performance event was caused by an issue on the cluster. You can identify the specific workload activities that might have been the main contributors to the performance event. IOPS are measured in operations per second (ops/sec).



This example shows the IOPS chart on the Performance/Volume Details page. The actual operations statistics is a blue line and the expected range of operations statistics is gray.



In some cases where a cluster is overloaded, Unified Manager might display the message Data collection is taking too long on Cluster cluster_name. This means that not enough statistics have been collected for Unified Manager to analyze. You need to reduce the resources the cluster is using so that statistics can be collected.

Performance monitoring of MetroCluster configurations

Unified Manager enables you to monitor the write throughput between clusters in a MetroCluster configuration to identify workloads with a high amount of write throughput. If these high-performing workloads are causing other volumes on the local cluster to have high I/O response times, Unified Manager triggers performance events to notify you.

When a local cluster in a MetroCluster configuration mirrors its data to its partner cluster, the data is written to NVRAM and then transferred over the interswitch links (ISLs) to the remote aggregates. Unified Manager analyzes the NVRAM to identify the workloads whose high write throughput is overutilizing the NVRAM, placing the NVRAM in contention.

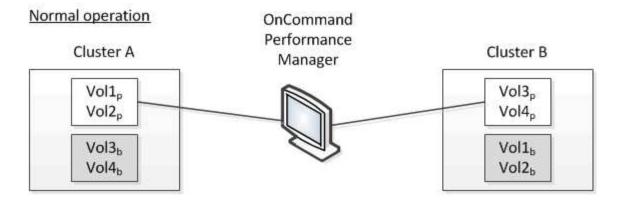
Workloads whose deviation in response time has exceeded the performance threshold are called *victims* and workloads whose deviation in write throughput to the NVRAM is higher than usual, causing the contention, are called *bullies*. Because only the write requests are mirrored to the partner cluster, Unified Manager does not analyze read throughput.

Unified Manager treats the clusters in a MetroCluster configuration as individual clusters. It does not distinguish between clusters that are partners or correlate the write throughput from each cluster.

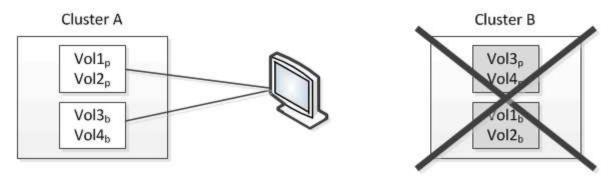
Volume behavior during switchover and switchback

Events that trigger a switchover or switchback cause active volumes to be moved from one cluster to the other cluster in the disaster recovery group. The volumes on the cluster that were active and serving data to clients are stopped, and the volumes on the other cluster are activated and start serving data. Unified Manager monitors only those volumes that are active and running.

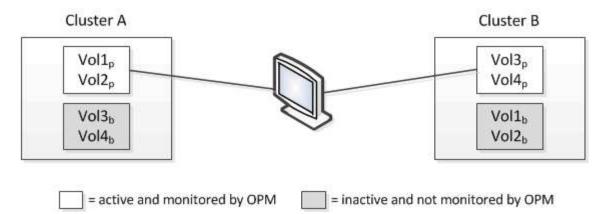
Because volumes are moved from one cluster to another, it is recommended that you monitor both clusters. A single instance of Unified Manager can monitor both clusters in a MetroCluster configuration, but sometimes the distance between the two locations necessitates using two Unified Manager instances to monitor both clusters. The following figure shows a single instance of Unified Manager:



Cluster B fails --- switchover to Cluster A



Cluster B is repaired --- switchback to Cluster B



The volumes with p in their names indicate the primary volumes, and the volumes with b in their names are mirrored backup volumes that are created by SnapMirror.

During normal operation:

- Cluster A has two active volumes: Vol1p and Vol2p.
- Cluster B has two active volumes: Vol3p and Vol4p.
- Cluster A has two inactive volumes: Vol3b and Vol4b.
- Cluster B has two inactive volumes: Vol1b and Vol2b.

Information pertaining to each of the active volumes (statistics, events, and so on) is collected by Unified Manager. Vol1p and Vol2p statistics are collected by Cluster A, and Vol3p and Vol4p statistics are collected by Cluster B.

After a catastrophic failure causes a switchover of active volumes from Cluster B to Cluster A:

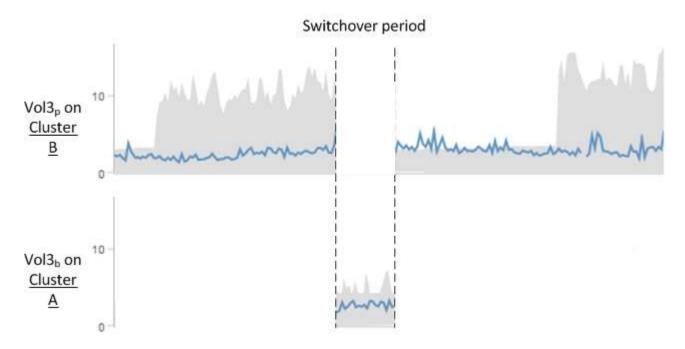
- Cluster A has four active volumes: Vol1p, Vol2p, Vol3b, and Vol4b.
- Cluster B has four inactive volumes: Vol3p, Vol4p, Vol1b, and Vol2b.

As during normal operation, information pertaining to each of the active volumes is collected by Unified Manager. But in this case, Vol1p and Vol2p statistics are collected by Cluster A, and Vol3b and Vol4b statistics are also collected by Cluster A.

Note that Vol3p and Vol3b are not the same volumes, because they are on different clusters. The information in Unified Manager for Vol3p is not the same as Vol3b:

- During switchover to Cluster A, Vol3p statistics and events are not visible.
- On the very first switchover, Vol3b looks like a new volume with no historical information.

When Cluster B is repaired and a switchback is performed, Vol3p is active again on Cluster B, with the historical statistics and a gap of statistics for the period during the switchover. Vol3b is not viewable from Cluster A until another switchover occurs:





- MetroCluster volumes that are inactive, for example, Vol3b on Cluster A after switchback, are identified with the message "This volume was deleted". The volume is not actually deleted, but it is not currently being monitored by Unified Manager because it is not the active volume.
- If a single Unified Manager is monitoring both clusters in a MetroCluster configuration, volume search returns information for whichever volume is active at that time. For example, a search for "Vol3" would return statistics and events for Vol3b on Cluster A if a switchover has occurred and Vol3 has become active on Cluster A.

What performance events are

Performance events are incidents related to workload performance on a cluster. They help you identify workloads with slow response times. Together with health events that

occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

When Unified Manager detects multiple occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events.

Performance event analysis and notification

Performance events notify you about I/O performance issues on a volume workload caused by contention on a cluster component. Unified Manager analyzes the event to identify all workloads involved, the component in contention, and whether the event is still an issue that you might need to resolve.

Unified Manager monitors the I/O latency (response time) and IOPS (operations) for volumes on a cluster. When other workloads overuse a cluster component, for example, the component is in contention and cannot perform at an optimal level to meet workload demands. The performance of other workloads that are using the same component might be impacted, causing their latencies to increase. If the latency crosses the performance threshold, Unified Manager triggers a performance event and sends an email alert to notify you.

Event analysis

Unified Manager performs the following analyses, using the previous 15 days of performance statistics, to identify the victim workloads, bully workloads, and the cluster component involved in an event:

- Identifies victim workloads whose latency has crossed the performance threshold, which is the upper boundary of the expected range:
 - For volumes on HDD or Flash Pool (hybrid) aggregates, events are triggered only when the latency is greater than 5 milliseconds (ms) and the IOPS are more than 10 operations per second (ops/sec).
 - For volumes on all-SSD aggregates or FabricPool (composite) aggregates, events are triggered only when the latency is greater than 1 ms and the IOPS are more than 100 ops/sec.
- Identifies the cluster component in contention.



If the latency of victim workloads at the cluster interconnect is greater than 1 ms, Unified Manager treats this as significant and triggers an event for the cluster interconnect.

- Identifies the bully workloads that are overusing the cluster component and causing it to be in contention.
- Ranks the workloads involved, based on their deviation in utilization or activity of a cluster component, to
 determine which bullies have the highest change in usage of the cluster component and which victims are
 the most impacted.

An event might occur for only a brief moment and then correct itself after the component it is using is no longer in contention. A continuous event is one that reoccurs for the same cluster component within a five-minute interval and remains in the active state. For continuous events, Unified Manager triggers an alert after detecting the same event during two consecutive analysis intervals. Events that remain unresolved, which have a state of new, can display different description messages as workloads involved in the event change.

When an event is resolved, it remains available in Unified Manager as part of the record of past performance issues for a volume. Each event has a unique ID that identifies the event type and the volumes, cluster, and cluster components involved.



A single volume can be involved in more than one event at the same time.

Event state

Events can be in one of the following states:

Active

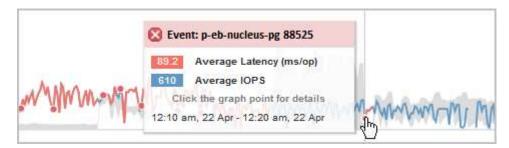
Indicates that the performance event is currently active (new or acknowledged). The issue causing the event has not corrected itself or has not been resolved. The performance counter for the storage object remains above the performance threshold.

Obsolete

Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the storage object is no longer above the performance threshold.

Event notification

The event alerts are displayed on the Dashboards/Overview page, Dashboards/Performance page, Performance/Volume Details page, and they are sent to specified email addresses. You can view detailed analysis information about an event and get suggestions for resolving it on the Event details page.



In this example, an event is indicated by a red dot () on the Latency chart on the Performance/Volume Details page. Hovering your mouse cursor over the red dot displays a popup with more details about the event and options for analyzing it.

Event interaction

On the Performance/Volume Details page, you can interact with events in the following ways:

- Moving the pointer over a red dot displays a message that shows the event ID, along with the latency, number of operations per second, and the date and time when the event was detected.
 - If there are multiple events for the same time period, the message shows the number of events, along with the average latency and operations per second for the volume.
- Clicking a single event displays a dialog box that shows more detailed information about the event, including the cluster components that are involved, similar to the Summary section on the Event details page.

The component in contention is circled and highlighted red. You can click either the event ID or **View full analysis** to view the full analysis on the Event details page. If there are multiple events for the same time period, the dialog box shows details about the three most recent events. You can click an event ID to view the event analysis on the Event details page. If there are more than three events for the same time period, clicking the red dot does not display the dialog box.

How Unified Manager determines the performance impact for an event

Unified Manager uses the deviation in activity, utilization, write throughput, cluster component usage, or I/O latency (response time) for a workload to determine the level of impact to workload performance. This information determines the role of each workload in the event and how they are ranked on the Event details page.

Unified Manager compares the last analyzed values for a workload to the expected range of values. The difference between the values last analyzed and the expected range of values identifies the workloads whose performance was most impacted by the event.

For example, suppose a cluster contains two workloads: Workload A and Workload B. The expected range for Workload A is 5-10 milliseconds per operation (ms/op) and its actual latency is usually around 7 ms/op. The expected range for Workload B is 10-20 ms/op and its actual latency is usually around 15 ms/op. Both workloads are well within their expected range for latency. Due to contention on the cluster, the latency of both workloads increases to 40 ms/op, crossing the performance threshold, which is the upper bounds of the expected range, and triggering events. The deviation in latency, from the expected values to the values above the performance threshold, for Workload A is around 33 ms/op, and the deviation for Workload B is around 25 ms/op. The latency of both workloads spike to 40 ms/op, but Workload A had the bigger performance impact because it had the higher latency deviation at 33 ms/op.

On the Event details page, in the System Diagnosis section, you can sort workloads by their deviation in activity, utilization, or throughput for a cluster component. You can also sort workloads by latency. When you select a sort option, Unified Manager analyzes the deviation in activity, utilization, throughput, or latency since the event was detected from the expected values to determine the workload sort order. For the latency, the red dots () indicate a performance threshold crossing by a victim workload, and the subsequent impact to the latency. Each red dot indicates a higher level of deviation in latency, which helps you identify the victim workloads whose latency was impacted the most by an event.

Cluster components and why they can be in contention

You can identify cluster performance issues when a cluster component goes into contention. The performance of volume workloads that use the component slow down and their response time (latency) for client requests increases, which triggers an event in Unified Manager.

A component that is in contention cannot perform at an optimal level. Its performance has declined, and the performance of other cluster components and workloads, called *victims*, might have increased latency. To bring a component out of contention, you must reduce its workload or increase its ability to handle more work, so that the performance can return to normal levels. Because Unified Manager collects and analyzes workload performance in five-minute intervals, it detects only when a cluster component is consistently overused. Transient spikes of overusage that last for only a short duration within the five-minute interval are not detected.

For example, a storage aggregate might be under contention because one or more workloads on it are competing for their I/O requests to be fulfilled. Other workloads on the aggregate can be impacted, causing their performance to decrease. To reduce the amount of activity on the aggregate, there are different steps you can take, such as moving one or more workloads to a less busy aggregate, to lessen the overall workload demand on the current aggregate. For a QoS policy group, you can adjust the throughput limit, or move workloads to a different policy group, so that the workloads are no longer being throttled.

Unified Manager monitors the following cluster components to alert you when they are in contention:

Network

Represents the wait time of I/O requests by the iSCSI protocols or the Fibre Channel (FC) protocols on the cluster. The wait time is time spent waiting for iSCSI Ready to Transfer (R2T) or FCP Transfer Ready (XFER_RDY) transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the block protocol layer is impacting the latency of one or more workloads.

Network Processing

Represents the software component in the cluster involved with I/O processing between the protocol layer and the cluster. The node handling network processing might have changed since the event was detected. If the network processing component is in contention, it means high utilization at the network processing node is impacting the latency of one or more workloads.

QoS Policy

Represents the storage Quality of Service (QoS) policy group of which the workload is a member. If the policy group component is in contention, it means all workloads in the policy group are being throttled by the set throughput limit, which is impacting the latency of one or more of those workloads.

Cluster Interconnect

Represents the cables and adapters with which clustered nodes are physically connected. If the cluster interconnect component is in contention, it means high wait time for I/O requests at the cluster interconnect is impacting the latency of one or more workloads.

Data Processing

Represents the software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains the workload. The node handling data processing might have changed since the event was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the latency of one or more workloads.

MetroCluster Resources

Represents the MetroCluster resources, including NVRAM and interswitch links (ISLs), used to mirror data between clusters in a MetroCluster configuration. If the MetroCluster component is in contention, it means high write throughput from workloads on the local cluster or a link health issue is impacting the latency of one or more workloads on the local cluster. If the cluster is not in a MetroCluster configuration, this icon is not displayed.

Aggregate or SSD Aggregate Ops

Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate). An "SSD Aggregate" consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

Cloud Latency

Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

Sync SnapMirror

Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

Roles of workloads involved in a performance event

Unified Manager uses roles to identify the involvement of a workload in a performance event. The roles include victims, bullies, and sharks. A user-defined workload can be a victim, bully, and shark at the same time.

Role	Description
Victim	A user-defined workload whose performance has decreased due to other workloads, called bullies, that are over-using a cluster component. Only user-defined workloads are identified as victims. Unified Manager identifies victim workloads based on their deviation in latency, where the actual latency, during an event, has greatly increased from its expected range of latency.
Bully	A user-defined or system-defined workload whose over-use of a cluster component has caused the performance of other workloads, called victims, to decrease. Unified Manager identifies bully workloads based on their deviation in usage of a cluster component, where the actual usage, during an event, has greatly increased from its expected range of usage.
Shark	A user-defined workload with the highest usage of a cluster component compared to all workloads involved in an event. Unified Manager identifies shark workloads based on their usage of a cluster component during an event.

Workloads on a cluster can share many of the cluster components, such as storage aggregates and the CPU for network and data processing. When a workload, such as a volume, increases its usage of a cluster component to the point that the component cannot efficiently meet workload demands, the component is in contention. The workload that is over-using a cluster component is a bully. The other workloads that share those components, and whose performance is impacted by the bully, are the victims. Activity from system-defined workloads, such as deduplication or Snapshot copies, can also escalate into "bullying".

When Unified Manager detects an event, it identifies all workloads and cluster components involved, including the bully workloads that caused the event, the cluster component that is in contention, and the victim workloads whose performance has decreased due to the increased activity of bully workloads.



If Unified Manager cannot identify the bully workloads, it only alerts on the victim workloads and the cluster component involved.

Unified Manager can identify workloads that are victims of bully workloads, and also identify when those same workloads become bully workloads. A workload can be a bully to itself. For example, a high-performing workload that is being throttled by a policy group limit causes all workloads in the policy group to be throttled, including itself. A workload that is a bully or a victim in an ongoing performance event might change its role or no longer be a participant in the event. On the Performance/Volume Details page, in the Events List table, when the selected volume changes its participant role, the date and time of the role change is displayed.

Analyzing workload performance

Unified Manager enables you to monitor and analyze I/O performance of volume workloads on your clusters. You can determine whether a performance issue is on the cluster and whether storage is the issue.



This chapter describes how to analyze workload performance using the Performance/Volume Details page and the Event details page.

Determining whether a workload has a performance issue

You can use Unified Manager to determine whether a detected performance event was truly caused by a performance issue on the cluster. The event might have been caused a spike in activity, for example, that drove up its response time, but now the response time has returned to it usual levels.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have identified the name of the volume, or associated LUN, you want to analyze.
- Unified Manager must have collected and analyzed a minimum of five days of performance statistics from the cluster.

About this task

If you are viewing the Event details page, you can click the name link for a volume to go directly to the Performance/Volume Details page.

Steps

1. In the **Search** bar, type at least the first three characters of the volume name.

The name of the volume is displayed in the search results.

2. Click the name of the volume.

The volume is displayed on the Performance/Volume Details page.

3. In the **Historic data** chart, click **5d** to display the last five days of historical data.

- 4. Review the **Latency** chart to answer the following questions:
 - Are there new performance events?
 - Are there historic performance events, indicating that the volume has had issues in the past?
 - · Are there spikes in the response time, even if the spikes are within the expected range?
 - On the Have there been configuration changes on the cluster that might have impacted performance? If the response time for the volume does not display performance events, spikes in activity, or recent configuration changes that might have impacted the response time, you can rule out the performance issue being caused by the cluster.

Investigating a perceived slow response time for a workload

You can use Unified Manager to determine whether operations on the cluster might have contributed to the slow response time (latency) for a volume workload.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have identified the name of the volume, or associated LUN, you want to analyze.
- Unified Manager must have collected and analyzed a minimum of five days of performance statistics from the cluster.

About this task

If you are viewing the Event details page, you can click the name for a volume to go directly to the Performance/Volume Details page.

Steps

1. In the **Search** bar, type the name of the volume.

The name of the volume is displayed in the search results.

2. Click the name of the volume.

The volume is displayed on the Performance/Volume Details page.

- 3. On the Historic data chart, click **5d** to display the last five days of historical data.
- 4. Review the **IOPS** chart to answer the following questions:
 - Are there dramatic spikes in the activity?
 - Are there dramatic drops in the activity?
 - Are there abnormal changes in the operations pattern? If the operations do not display dramatic spikes or drops in activity, and there were no changes to the cluster configuration during this time, the storage administrator can confirm that other workloads have not impacted volume performance.
- 5. On the Break down data by menu, under IOPS, select Reads/writes/other.
- 6. Click Submit.

The Reads/writes/other chart is displayed below the IOPS chart.

Review the Reads/writes/other chart to identify dramatic spikes or drops in the amount of reads or writes for the volume.

If there are no dramatic spikes or drops in reads or writes, the storage administrator can confirm that I/O on the cluster is operating normally. Any performance issues might be on the network or the connected clients.

Identifying trends of I/O response time on cluster components

You can use Unified Manager to view the performance trends for all monitored cluster components for a volume workload. You can see, over time, which components have the highest usage, whether the highest usage is from read or write requests, and how the usage has impacted the workload response time.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have identified the name of the volume or associated LUN you want to analyze.
- To display 30 days of performance statistics, Unified Manager must have collected and analyzed a minimum of 30 days of performance statistics from the cluster.

About this task

Identifying performance trends for the cluster components helps the administrator decide whether the cluster is being overused or underused.

If you are viewing the Event details page, you can click the name for a volume to go directly to the Performance/Volume Details page.

Steps

1. In the **Search** bar, type the name of the volume.

The name of the volume is displayed in the search results.

2. Click the name of the volume.

The volume is displayed on the Performance/Volume Details page.

- 3. On the Historic data chart, click **30d** to display the last 30 days of historical data.
- 4. Click Break down data by.
- 5. Under Latency, select Cluster Components and Reads/writes latency.
- 6. Click Submit.

Both charts are displayed below the Latency chart.

7. Review the **Cluster Components** chart.

The chart breaks down the total response time by cluster component. The response time at the aggregate is the highest.

8. Compare the Cluster Components chart to the Latency chart.

The Latency chart shows spikes in the total response time that are aligned with the spikes in response time for the aggregate. There are a few at the end of the 30-day time frame, where the performance threshold was crossed.

9. Review the Reads/writes latency chart.

The chart shows a higher response time for write requests than read requests, indicating that the client applications are waiting longer than usual to have their write requests fulfilled.

10. Compare the Reads/writes latency chart to the Latency chart.

The spikes in total response time that align with the aggregate in the Cluster Components chart also align with the writes in the Reads/writes latency chart. The administrator must decide whether the client applications using the workload must be addressed or whether the aggregate is being overused.

Analyzing the performance improvements achieved from moving a volume

You can use Unified Manager to investigate the impact of a volume move operation on the latency (response time) of other volumes on the cluster. Moving a high performing volume to a less busy aggregate or an aggregate with flash storage enabled allows the volume to perform more efficiently.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have identified the name of the volume, or associated LUN, you want to analyze.
- Unified Manager must have collected and analyzed seven days of data.

About this task

Unified Manager identifies when a volume moves between aggregates. It can detect when the volume move is occurring, completed, or failed. The Performance/Volume Details page displays a change event icon for each state of the volume move, which helps you track when a move operation occurred and helps you determine whether it might have contributed to a performance event.

If you are viewing the Event details page, you can click the name of a volume to go directly to the Performance/Volume Details page.

Steps

- 1. In the **Search** bar, type the name of the volume.
- 2. Click the name of the volume.

The volume is displayed on the Performance/Volume Details page.

- 3. In the **Historic data** chart, adjust the sliders to show activity from the previous work week.
- 4. Analyze the **Latency** chart and the **IOPS** chart to see how the volume performed over the last few days.

Assume that you notice a consistent pattern of very high average response times of over 42 milliseconds per operation (ms/op), with performance events, each day of the week and decide to move the volume to a less busy aggregate to improve performance. Using OnCommand System Manager, you can move the

volume to an aggregate with Flash Pool enabled for an increased performance boost. Approximately an hour after the volume move has been completed, you can return to Unified Manager to confirm that the move operation was completed successfully and that the latency has improved.

- 5. If the **Performance/Volume Details** page is not displayed, search for the volume you want to view.
- 6. On the **Historic data** chart, click **1d** to view the activity from the last one day, a few hours since the volume move was completed.

At the bottom of the page, in the Events time line, a change event icon () is displayed to indicate the time that the volume move operation was completed. A black, vertical line is also displayed from the change event icon to the Latency chart.

7. Point your cursor to the change event icon to view details about the event in the Events List.

Because the volume moved to an aggregate with Flash Pool enabled, you can see the change in read and write I/O to cache.

8. On the Break down data by menu, under MBps, select Cache hit ratio.

The Cache hit ratio chart displays statistics about the reads and writes to cache.

The volume successfully moved to a less busy aggregate and the change event is highlighted in the Events List on the right. The average latency decreased significantly from over 42 ms/op to around 24 ms/op. The current latency is around 1.5 ms/op. In the Cache hit ratio chart, the amount of successful read and write hits to cache is now at 100% because the volume is now on an aggregate with Flash Pool enabled.

How moving a FlexVol volume works

Knowing how moving a FlexVol volume works helps you to determine whether the volume move satisfies service-level agreements and to understand where a volume move is in the volume move process.

FlexVol volumes are moved from one aggregate or node to another within the same storage virtual machine (SVM). A volume move does not disrupt client access during the move.

Moving a volume occurs in multiple phases:

- A new volume is made on the destination aggregate.
- The data from the original volume is copied to the new volume.

During this time, the original volume is intact and available for clients to access.

At the end of the move process, client access is temporarily blocked.

During this time the system performs a final replication from the source volume to the destination volume, swaps the identities of the source and destination volumes, and changes the destination volume to the source volume.

 After completing the move, the system routes client traffic to the new source volume and resumes client access.

The move is not disruptive to client access because the time in which client access is blocked ends before

clients notice a disruption and time out. Client access is blocked for 35 seconds by default. If the volume move operation cannot finish in the time that access is denied, the system aborts this final phase of the volume move operation and allows client access. The system attempts the final phase three times by default. After the third attempt, the system waits an hour before attempting the final phase sequence again. The system runs the final phase of the volume move operation until the volume move is complete.

Performance/Volume Details page

This page provides detailed performance statistics for all I/O activity and operations for the selected FlexVol volume, FlexGroup volume, or FlexGroup constituent workload. You can select a specific time frame over which to view the statistics and events for the volume. The events identify performance events and changes that might be impacting I/O performance.

Historic data chart

Plots the historical performance analysis data for the selected volume. You can click and drag the sliders to specify a time frame. The sliders increase and decrease the time frame window. The data outside the time frame window is grayed out. You can use the slider at the bottom of the chart to move the time frame window across the historical data. The entire page, including the displayed charts and events, reflects the data available within the time frame window. Unified Manager retains a maximum of 30 days of historical data on this page.



On the historic data chart, if you select a time frame of more than 1 day, depending on your screen resolution, the charts display the maximum values for response time and IOPS across the number of days.

Options

Time selector

Specifies the time range over which to view the volume performance statistics for the entire page. You can click 1 day (1d) through 30 days (30d), or click **Custom** to select a custom range. For a custom range, you can select a beginning and end date, and then click **Update** to update the entire page.



If you access the Performance/Volume Details page by clicking the name link of a volume on the Event details page, a time range, such as 1 day or 5 days prior to the current day, is automatically selected by default. When you move the slider in the historic data chart, the time range changes to a custom range, but the **Custom** time selector is not selected. The default time selector remains selected.

Break down data by

Provides a list of charts you can add to the Performance/Volume Details page to display more detailed performance statistics for the selected volume.

Performance statistics displayed in the data breakdown charts

You can use the graphs to view performance trending for a volume. You can also view statistics for reads and writes, network protocol activity, the impact of QoS policy group throttling on latency, the ratio of reads and writes to cache storage, the total cluster CPU

time used by a workload, and specific cluster components.

These views display a maximum of 30 days of statistics from the current day. On the historic data chart, if you select a time frame of more than 1 day, depending on your screen resolution, the charts display the maximum values for latency and IOPS across the number of days.



You can use the Select All check box to select, or deselect, all the listed chart options.

Latency

The following charts detail the latency, or response time, information for the selected workload:

Cluster Components

Displays a graph of the time spent at each cluster component used by the selected volume.

The chart helps you determine the latency impact by each component as it relates to the total latency. You can use the check box next to each component to show and hide its graph.

For QoS policy groups, data is only displayed for user-defined policy groups. Zeros are displayed for system-defined policy groups, such as default policy groups.

Reads/writes latency

Displays a graph of the latencies of the successful read and write requests from the selected volume workload over the selected time frame.

Write requests are an orange line and read requests are a blue line. The requests are specific to the latency for the selected volume workload, not all workloads on the cluster.



The read and write statistics might not always add up to the total latency statistics displayed in the Latency chart. This is expected behavior based on how Unified Manager collects and analyzes read and write statistics for a workload.

Policy Group Impact

Displays a graph of the percentage of the latency for the selected volume workload that is impacted by the throughput limit on its QoS policy group.

If the workload is throttled, the percentage indicates how much the throttling contributed to the latency at a specific point in time. The percentage values indicate the amount of throttling:

- 0% = no throttling
- > 0% = throttling
- > 20% = critical throttling
 If the cluster can handle more work, you can reduce throttling by increasing the policy group limit.
 Another option is to move the workload to a less busy aggregate.



The chart displays for workloads in a user-defined QoS policy group with a set throughput limit only. It does not display if the workloads are in a system-defined policy group, such as the default policy group, or a policy group that does not have a QoS limit. For a QoS policy group, you can point the cursor to the name of the policy group to display its throughput limit and the last time it was modified. If the policy group was modified before the associated cluster was added to Unified Manager, the last modified time is the date and time when Unified Manager first discovered the cluster.

· IOPS

The following charts detail the IOPS data for the selected workload:

Reads/writes/other

Displays a graph showing the number of read and write IOPS and other IOPS, per second, over the selected time frame.

Other IOPS are protocol activities initiated by the client that are not reads or writes. For example, in an NFS environment, this could be metadata operations such as getattr, setattr, or fsstat. In a CIFS environment, this could be attribute lookups, directory listings, or antivirus scans. Write IOPS are an orange line and read requests are a blue line. The requests are specific to all operations for the selected volume workload, not all operations on the cluster.

MBps

The following charts detail the throughput data for the selected workload:

Cache hit ratio

Displays a graph of the percentage of read requests from client applications satisfied by cache over the selected time frame.

The cache could be on Flash Cache cards or solid state drives (SSDs) in Flash Pool aggregates. A cache hit, in blue, is a read from cache. A cache miss, in orange, is a read from a disk in the aggregate. The requests are specific to the selected volume workload, not all workloads on the cluster.

You can view more detailed information about volume cache usage in the Unified Manager Health pages and in OnCommand System Manager.

Components

The following charts detail the data by cluster component used by the selected workload:

Cluster CPU Time

Displays a graph of the CPU usage time, in ms, for all nodes in the cluster used by the selected workload.

The graph displays the combined CPU usage time for network processing and data processing. The CPU time for system-defined workloads that are associated to the selected workload, and are using the same nodes for data processing, is also included. You can use the chart to determine whether the workload is a high consumer of the CPU resources on the cluster. You can also use the chart, in combination with the Reads/writes latency chart under the Latency chart, or the Reads/writes/other chart under the IOPS chart, to determine how changes to workload activity over time impact cluster CPU utilization.

Disk Utilization

Displays a graph showing the percentage of utilization on the data disks in the storage aggregate over the selected time frame

The utilization includes disk read and write requests from the selected volume workload only. Reads from cache are not included. The utilization is specific to the selected volume workload, not all workloads on the disks. If a monitored volume is involved in a volume move, the utilization values in this chart are for the target aggregate to which the volume moved.

How graphs of performance data work

Unified Manager uses graphs or charts to show you volume performance statistics and events over a specified period of time.

The graphs enable you to customize the range of time for which to view data. The data is displayed with the time frame on the horizontal axis of the graph and the counters on the vertical axis, with point intervals along the graph lines. The vertical axis is dynamic; the values adjust based on the peaks of the expected or actual values.

Selecting time frames

On the Performance/Volume Details page, the Historic data chart enables you to select a time frame for all graphs on the page. The 1d, 5d, 10d, and 30d buttons specify 1 day through 30 days (1 month) and the **Custom** button enables you to specify a custom time range within that 30 days. Each point on a graph represents a 5-minute collection interval, and a maximum of 30 days of historical performance data is retained. Note that intervals also account for network delays and other anomalies.



In this example, the Historic data chart has a time frame set to the beginning and the end of the month of March. In the selected time frame, all historic data before March is grayed out.

Viewing data point information

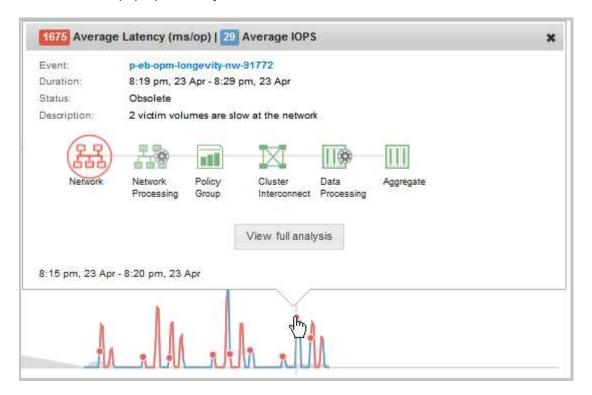
To view data point information on a graph, you can position the cursor over a specific point within the graph, and a pop-up box displays listing the value and date and time information.



In this example, positioning the cursor over the IOPS chart on the Performance/Volume Details page displays the response time and operations values between 3:50 a.m. and 3:55 a.m. on October 20th.

Viewing performance event information

To view event information on a graph, you can position the cursor over an event icon to view summary information in a pop-up box, or you can click the event icon for more detailed information.



In this example, on the Performance/Volume Details page, clicking an event icon on the Latency chart displays detailed information about the event in a pop-up box. The event is also highlighted in the Events List.

Analyzing performance events

You can analyze performance events to identify when they were detected, whether they are active (new or acknowledged) or obsolete, the workloads and cluster components involved, and the options for resolving the events on your own.

Displaying information about performance events

You can use the Events inventory page to view a list of all the new and obsolete performance events on the clusters being monitored by Unified Manager. By viewing this information you can determine the most critical events and then drill down to detailed information to determine the cause of the event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

About this task

The list of events is sorted by detected time, with the most recent events listed first. You can click a column header to sort the events based on that column. For example, you can sort by the Status column to view

events by severity. If you are looking for a specific event, or for a specific type of event, you can use the filter and search mechanisms to refine the list of events that appear in the list.

Events from all sources are displayed on this page:

- User-defined performance threshold policy
- System-defined performance threshold policy
- · Dynamic performance threshold

The Event Type column lists the source of the event. You can select an event to view details about the event in the Event details page.

Steps

- 1. In the left navigation pane, click **Events**.
- 2. Locate an event that you want to analyze and click the event name.

The details page for the event displays.



You can also display the details page for an event by clicking the event name link from the Performance Explorer page and from an alert email.

Analyzing events from user-defined performance thresholds

Events generated from user-defined thresholds indicate that a performance counter for a certain storage object, for example, an aggregate or volume, has crossed the threshold you defined in the policy. This indicates that the cluster object is experiencing a performance issue.

You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.

Responding to user-defined performance threshold events

You can use Unified Manager to investigate performance events caused by a performance counter crossing a user-defined warning or critical threshold. You can also use Unified Manager to check the health of the cluster component to see whether recent health events detected on the component contributed to the performance event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

Steps

- 1. Display the **Event** details page to view information about the event.
- 2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message "Latency value of 456 ms/op has triggered a WARNING event based on

threshold setting of 400 ms/op" indicates that a latency warning event occurred for the object.

3. Hover your cursor over the policy name to display details about the threshold policy that triggered the event.

This includes the policy name, the performance counter being evaluated, the counter value that must be breached to be considered a critical or warning event, and the duration by which the counter must exceed the value.

- 4. Make a note of the **Event Trigger Time** so you can investigate whether other events might have occurred at the same time that could have contributed to this event.
- 5. Follow one of the options below to further investigate the event, to determine whether you need to perform any actions to resolve the performance problem:

Option	Possible investigation actions
Click the Source object name to display the Explorer page for that object.	This page enables you to view the object details and compare this object with other similar storage objects to see whether other storage objects have a performance issue around the same time. For example, to see whether other volumes on the same aggregate are also having a performance issue.
Click the cluster name to display the Cluster Summary page.	This page enables you to view the details for the cluster on which this object resides to see whether other performance issues have occurred around the same time.

Analyzing events from system-defined performance thresholds

Events generated from system-defined performance thresholds indicate that a performance counter, or set of performance counters, for a certain storage object has crossed the threshold from a system-defined policy. This indicates that the storage object, for example, an aggregate or node, is experiencing a performance issue.

You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.



System-defined threshold policies are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

Responding to system-defined performance threshold events

You can use Unified Manager to investigate performance events caused by a performance counter crossing a system-defined warning threshold. You can also use Unified Manager to check the health of the cluster component to see whether recent events detected on the component contributed to the performance event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

Steps

- 1. Display the **Event** details page to view information about the event.
- 2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message "Node utilization value of 90 % has triggered a WARNING event based on threshold setting of 85 %" indicates that a node utilization warning event occurred for the cluster object.

- 3. Make a note of the **Event Trigger Time** so you can investigate whether other events might have occurred at the same time that could have contributed to this event.
- 4. Under **System Diagnosis**, review the brief description of the type of analysis the system-defined policy is performing on the cluster object.

For some events a green or red icon is displayed next to the diagnosis to indicate whether an issue was found in that particular diagnosis. For other types of system-defined events counter charts display the performance for the object.

5. Under **Suggested Actions**, click the **Help me do this** link to view the suggested actions you can perform to try and resolve the performance event on your own.

Responding to QoS policy group performance events

Unified Manager generates QoS policy warning events when workload throughput (IOPS, IOPS/TB, or MBps) has exceeded the defined ONTAP QoS policy setting and workload latency is becoming affected. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

About this task

Unified Manager generates warning events for QoS policy breaches when workload throughput has exceeded the defined QoS policy setting during each performance collection period for the previous hour. Workload throughput may exceed the QoS threshold for only a short period of time during each collection period, but Unified Manager displays only the "average" throughput during the collection period on the chart. For this reason you may receive QoS events while the throughput for a workload might not have crossed the policy threshold shown in the chart.

You can use System Manager or the ONTAP commands to manage policy groups, including the following tasks:

- Creating a new policy group for the workload
- Adding or removing workloads in a policy group

- Moving a workload between policy groups
- · Changing the throughput limit of a policy group
- · Moving a workload to a different aggregate or node

Steps

- 1. Display the **Event** details page to view information about the event.
- 2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message "IOPS value of 1,352 IOPS on vol1_NFS1 has triggered a WARNING event to identify potential performance problems for the workload" indicates that a QoS Max IOPS event occurred on volume vol1_NFS1.

3. Review the **Event Information** section to see more details about when the event occurred and how long the event has been active.

Additionally, for volumes or LUNs that are sharing the throughput of a QoS policy you can see the names of the top three workloads that are consuming the most IOPS or MBps.

4. Under the System Diagnosis section, review the two charts: one for total average IOPS or MBps (depending on the event), and one for latency. When arranged this way you can see which cluster components are most affecting latency when the workload approached the QoS max limit.

For a shared QoS policy event, the top three workloads are shown in the throughput chart. If more than three workloads are sharing the QoS policy, then additional workloads are added together in an "Other workloads" category. Additionally, the Latency chart shows the average latency on all workloads that are part of the QoS policy.

Note that for adaptive QoS policy events that the IOPS and MBps charts show IOPS or MBps values that ONTAP has converted from the assigned IOPS/TB threshold policy based on the size of the volume.

5. Under the **Suggested Actions** section, review the suggestions and determine which actions you should perform to avoid an increase in latency for the workload.

If required, click the **Help** button to view more details about the suggested actions you can perform to try and resolve the performance event.

Understanding events from adaptive QoS policies that have a defined block size

Adaptive QoS policy groups automatically scale a throughput ceiling or floor based on the volume size, maintaining the ratio of IOPS to TBs as the size of the volume changes. Starting with ONTAP 9.5 you can specify the block size in the QoS policy to effectively apply a MBps threshold at the same time.

Assigning an IOPS threshold in an adaptive QoS policy places a limit only on the number of operations that occur in each workload. Depending on the block size that is set on the client that generates the workloads, some IOPS include much more data and therefore place a much larger burden on the nodes that process the operations.

The MBps value for a workload is generated using the following formula:

```
MBps = (IOPS * Block Size) / 1000
```

If a workload is averaging 3,000 IOPS and the block size on the client is set to 32 KB, then the effective MBps for this workload is 96. If this same workload is averaging 3,000 IOPS and the block size on the client is set to 48 KB, then the effective MBps for this workload is 144. You can see that the node is processing 50% more data when the block size is larger.

Let's look at the following adaptive QoS policy that has a defined block size and how events are triggered based on the block size that is set on the client.

Create a policy and set the peak throughput to 2,500 IOPS/TB with a block size of 32KB. This effectively sets the MBps threshold to 80 MBps ((2500 IOPS * 32KB) / 1000) for a volume with 1 TB used capacity. Note that Unified Manager generates a Warning event when the throughput value is 10% less than the defined threshold. Events are generated under the following situations:

Used Capacity	Event is generated when throughput exceeds this number of
IOPS	MBps
1 TB	2,250 IOPS
72 MBps	2 TB
4,500 IOPS	144 MBps
5 TB	11,250 IOPS

If the volume is using 2TB of the available space, and the IOPS is 4,000, and the QoS block size is set to 32KB on the client, then the MBps throughput is 128 MBps ((4,000 IOPS * 32 KB) / 1000). No event is generated in this scenario because both 4,000 IOPS and 128 MBps are below the threshold for a volume that is using 2 TB of space.

If the volume is using 2TB of the available space, and the IOPS is 4,000, and the QoS block size is set to 64KB on the client, then the MBps throughput is 256 MBps ((4,000 IOPS * 64 KB) / 1000). In this case the 4,000 IOPS does not generate an event, but the MBps value of 256 MBps is above the threshold of 144 MBps and an event is generated.

For this reason, when an event is triggered based on a MBps breach for an adaptive QoS policy that includes the block size, a MBps chart is displayed in the System Diagnosis section of the Event details page. If the event is triggered based on an IOPS breach for the adaptive QoS policy, an IOPS chart is displayed in the System Diagnosis section. If a breach occurs for both IOPS and MBps you will receive two events.

For more information on adjusting QoS settings, see the ONTAP 9 Performance Monitoring Power Guide.

ONTAP 9 Performance Monitoring Power Guide

Responding to node resources overutilized performance events

Unified Manager generates node resources overutilized warning events when a single node is operating above the bounds of its operational efficiency, and therefore potentially affecting workload latencies. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

About this task

Unified Manager generates warning events for node resources overutilized policy breaches by looking for nodes that are using more than 100% of their performance capacity for more than 30 minutes.

You can use System Manager or the ONTAP commands to correct this type of performance issue, including the following tasks:

- Creating and applying a QoS policy to any volumes or LUNs that are overusing system resources
- Reducing the QoS maximum throughput limit of a policy group to which workloads have been applied
- · Moving a workload to a different aggregate or node
- Increasing capacity by adding disks to the node, or by upgrading to a node with a faster CPU and more RAM

Steps

- 1. Display the **Event** details page to view information about the event.
- 2. Review the **Description**, which describes the threshold breach that caused the event.

For example, the message "Perf. Capacity Used value of 139% on simplicity-02 has triggered a WARNING event to identify potential performance problems in the data processing unit." indicates that performance capacity on node simplicity-02 is overused and affecting node performance.

3. Under the **System Diagnosis** section, review the three charts: one for performance capacity used on the node, one for average storage IOPS being used by the top workloads, and one for latency on the top workloads. When arranged in this way you can see which workloads are the cause of the latency on the node.

You can view which workloads have QoS policies applied, and which do not, by moving your cursor over the IOPS chart.

4. Under the **Suggested Actions** section, review the suggestions and determine which actions you should perform to avoid an increase in latency for the workload.

If required, click the **Help** button to view more details about the suggested actions you can perform to try and resolve the performance event.

Analyzing events from dynamic performance thresholds

Events generated from dynamic thresholds indicate that the actual response time (latency) for a workload is too high, or too low, compared to the expected response time range. You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.



Dynamic performance thresholds are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

Identifying victim workloads involved in a dynamic performance event

In Unified Manager, you can identify which volume workloads have the highest deviation in response time (latency) caused by a storage component in contention. Identifying these workloads helps you understand why the client applications accessing them have been performing slower than usual.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete dynamic performance events.

About this task

The Event details page displays a list of the user-defined and system-defined workloads, ranked by the highest deviation in activity or usage on the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. In the Workload Latency and Workload Activity charts, select Victim Workloads.
- 3. Hover your cursor over the charts to view the top user-defined workloads that are affecting the component, and the name of the victim workload.

Identifying bully workloads involved in a dynamic performance event

In Unified Manager, you can identify which workloads have the highest deviation in usage for a cluster component in contention. Identifying these workloads helps you understand why certain volumes on the cluster have slow response times (latency).

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete dynamic performance events.

About this task

The Event details page displays a list of the user-defined and system-defined workloads ranked by the highest usage of the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. In the Workload Latency and Workload Activity charts, select Bully Workloads.
- 3. Hover your cursor over the charts to view the top user-defined bully workloads that are affecting the

Identifying shark workloads involved in a dynamic performance event

In Unified Manager, you can identify which workloads have the highest deviation in usage for a storage component in contention. Identifying these workloads helps you determine if these workloads should be moved to a less-utilized cluster.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- · There are new, acknowledged, or obsolete performance dynamic event.

About this task

The Event details page displays a list of the user-defined and system-defined workloads ranked by the highest usage of the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. In the Workload Latency and Workload Activity charts, select **Shark Workloads**.
- 3. Hover your cursor over the charts to view the top user-defined workloads that are affecting the component, and the name of the shark workload.

Performance event analysis for a MetroCluster configuration

You can use Unified Manager to analyze a performance event for a MetroCluster configuration. You can identify the workloads involved in the event and review the suggested actions for resolving it.

MetroCluster performance events might be due to *bully* workloads that are over-utilizing the interswitch links (ISLs) between the clusters, or due to link health issues. Unified Manager monitors each cluster in a MetroCluster configuration independently, without consideration of performance events on a partner cluster.

Performance events from both clusters in the MetroCluster configuration are also displayed on the Unified ManagerDashboards/Overview page. You can also view the Health pages of Unified Manager to check the health of each cluster and to view their relationship.

Analyzing a dynamic performance event on a cluster in a MetroCluster configuration

You can use Unified Manager to analyze the cluster in a MetroCluster configuration on which a performance event was detected. You can identify the cluster name, event detection time, and the *bully* and *victim* workloads involved.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events for a MetroCluster configuration.

• Both clusters in the MetroCluster configuration must be monitored by the same instance of Unified Manager.

Steps

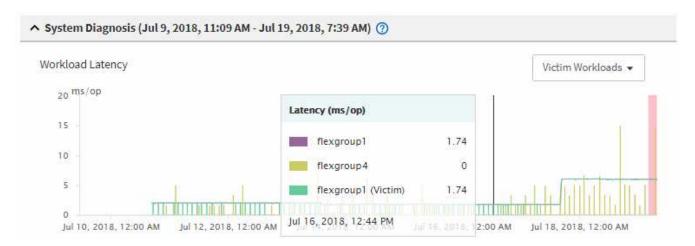
- 1. Display the **Event details** page to view information about the event.
- 2. Review the event description to see the names of the workloads involved and the number of workloads involved.

In this example, the MetroCluster Resources icon is red, indicating that the MetroCluster resources are in contention. You position your cursor over the icon to display a description of the icon. At the top of the page in the event ID, the cluster name identifies the name of the cluster on which the event was detected.



- 3. Make a note of the cluster name and the event detection time, which you can use to analyze performance events on the partner cluster.
- 4. In the charts, review the *victim* workloads to confirm that their response times are higher than the performance threshold.

In this example, the victim workload is displayed in the hover text. The Latency charts display, at a high-level, a consistent latency pattern for the victim workloads involved. Even though the abnormal latency of the victim workloads triggered the event, a consistent latency pattern might indicate that the workloads are performing within their expected range, but that a spike in I/O increased the latency and triggered the event.



If you recently installed an application on a client that accesses these volume workloads and that application sends a high amount of I/O to them, you might be anticipating their latencies to increase. If the latency for the workloads returns within the expected range, the event state changes to obsolete, and remains in this state for more than 30 minutes, you can probably ignore the event. If the event is ongoing, and remains in the new state, you can investigate it further to determine whether other issues caused the event.

In the Workload Throughput chart, select Bully Workloads to display the bully workloads.

The presence of bully workloads indicates that the event might have been caused by one or more workloads on the local cluster overutilizing the MetroCluster resources. The bully workloads have a high deviation in write throughput (MBps).

This chart displays, at a high-level, the write throughput (MBps) pattern for the workloads. You can review the write MBps pattern to identify abnormal throughput, which might indicate that a workload is over-utilizing the MetroCluster resources.

If no bully workloads are involved in the event, the event might have been caused by a health issue with the link between the clusters or a performance issue on the partner cluster. You can use Unified Manager to check the health of both clusters in a MetroCluster configuration. You can also use Unified Manager to check for and analyze performance events on the partner cluster.

Analyzing a dynamic performance event for a remote cluster on a MetroCluster configuration

You can use Unified Manager to analyze dynamic performance events on a remote cluster in a MetroCluster configuration. The analysis helps you determine whether an event on the remote cluster caused an event on its partner cluster.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have analyzed a performance event on a local cluster in a MetroCluster configuration and obtained the event detection time.
- You must have checked the health of the local cluster and its partner cluster involved in the performance event and obtained the name of the partner cluster.

Steps

- 1. Log in to the Unified Manager instance that is monitoring the partner cluster.
- 2. In the left navigation pane, click **Events** to display the event list.
- 3. From the Time Range selector, select Last Hour, and then click Apply Range.
- 4. In the **Filtering** selector, select **Cluster** from the left drop-down menu, type the name of the partner cluster in the text field, and then click **Apply Filter**.

If there are no events for the selected cluster over the last hour, this indicates that the cluster has not experienced any performance issues during the time that the event was detected on its partner.

5. If the selected cluster has events detected over the last hour, compare the event detection time to the event detection time for the event on the local cluster.

If these events involve bully workloads causing contention on the data processing component, one or more of these bullies might have caused the event on the local cluster. You can click the event to analyze it and review the suggested actions for resolving it on the Event details page.

If these events do not involve bully workloads, they did not cause the performance event on the local cluster.

Responding to a dynamic performance event caused by QoS policy group throttling

You can use Unified Manager to investigate a performance event caused by a Quality of Service (QoS) policy group throttling workload throughput (MBps). The throttling increased the response times (latency) of volume workloads in the policy group. You can use the event information to determine whether new limits on the policy groups are needed to stop the throttling.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. Read the **Description**, which displays the name of the workloads impacted by the throttling.



The description can display the same workload for the victim and bully, because the throttling makes the workload a victim of itself.

3. Record the name of the volume, using an application such as a text editor.

You can search on the volume name to locate it later.

- 4. In the Workload Latency and Workload Activity charts, select **Bully Workloads**.
- 5. Hover your cursor over the charts to view the top user-defined workloads that are affecting the policy group.

The 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.

- 6. Navigate to the **Performance/Volume Details** page for the top workload.
- Select Break down data by.
- 8. Select the check box next to **Latency** to select all latency breakdown charts.
- 9. Under IOPS, select Reads/writes/other.
- 10. Click Submit.

The breakdown charts are displayed under the Latency chart and the IOPS chart.

11. Compare the **Policy Group Impact** chart to the **Latency** chart to see what percentage of throttling impacted the latency at the time of the event.

The policy group has a maximum throughput of 1,000 operations per second (op/sec), which the workloads in it cannot collectively exceed. At the time of the event, the workloads in the policy group had a combined throughput of over 1,200 op/sec, which caused the policy group to throttle its activity back to 1,000 op/sec. The Policy Group Impact chart shows that the throttling caused 10% of the total latency, confirming that the throttling caused the event to occur.

12. Review the **Cluster Components** chart, which shows the total latency by cluster component.

The latency is highest at the policy group, further confirming that the throttling caused the event.

13. Compare the Reads/writes latency 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 there is a high amount of throughput or number of operations that increased the latency. You can use these values when deciding to put a policy group limit on the throughput or operations.

- 14. Use OnCommand System Manager to increase the current limit on the policy group to 1,300 op/sec.
- 15. After a day, return to Unified Manager and search for the name of the workload that you recorded in Step 3.

The Performance/Volume Details page is displayed.

- 16. Select Break down data by > IOPS.
- 17. Click Submit.

The Reads/writes/other chart is displayed.

- 18. At the bottom of the page, point your cursor to the change event icon () for the policy group limit change.
- 19. Compare the **Reads/writes/other** chart to the **Latency** chart.

The read and write requests are the same, but the throttling has stopped and the latency has decreased.

Responding to a dynamic performance event caused by a disk failure

You can use Unified Manager to investigate a performance event caused by workloads overutilizing an aggregate. You can also use Unified Manager to check the health of the aggregate to see if recent health events detected on the aggregate contributed to the performance event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

- 1. Display the **Event details** page to view information about the event.
- 2. Read the **Description**, which describes the workloads involved in the event and the cluster component in contention

There are multiple victim volumes whose latency was impacted by the cluster component in contention. The aggregate, which is in the middle of a RAID reconstruct to replace the failed disk with a spare disk, is the cluster component in contention. Under Component in Contention, the Aggregate icon is highlighted red and the name of the aggregate is displayed in parentheses.

- In the Workload Utilization chart, select Bully Workloads.
- 4. Hover your cursor over the chart to view the top bully workloads that are affecting the component.

The top workloads with the highest peak utilization since the event was detected are displayed at the top of the chart. One of the top workloads is the system-defined workload Disk Health, which indicates a RAID reconstruct. A reconstruct is the internal process involved with rebuilding the aggregate with the spare disk. The Disk Health workload, along with other workloads on the aggregate, likely caused the contention on the aggregate and the associated event.

- After confirming that the activity from the Disk Health workload caused the event, wait for approximately 30
 minutes for the reconstruction to finish and for Unified Manager to analyze the event and detect whether
 the aggregate is still in contention.
- 6. In Unified Manager, search for the event ID you recorded in Step 2.

The event for the disk failure is displayed on the Event details page. After the RAID reconstruction is complete, check that the State is obsolete, indicating that the event is resolved.

- 7. In the Workload Utilization chart, select **Bully Workloads** to view the workloads on the aggregate by peak utilization.
- 8. Navigate to the **Performance/Volume Details** page for the top workload.
- 9. Click 1d to display the last 24 hours (1 day) of data for the selected volume.

In the Latency chart, a red dot () indicates when the disk failure event occurred.

- 10. Select Break down data by.
- 11. Under Components, select Disk Utilization.
- 12. Click Submit.

The Disk Utilization chart displays a graph of all read and write requests from the selected workload to the disks of the target aggregate.

13. Compare the data in the **Disk Utilization** chart to the data at the time of the event in the **Latency** chart.

At the time of the event, the Disk Utilization shows a high amount of read and write activity, caused by the RAID reconstruction processes, which increased the latency of the selected volume. A few hours after the event occurred, both the reads and writes and the latency have decreased, confirming that the aggregate is no longer in contention.

Responding to a dynamic performance event caused by HA takeover

You can use Unified Manager to investigate a performance event caused by high data processing on a cluster node that is in a high-availability (HA) pair. You can also use Unified Manager to check the health of the nodes to see whether any recent health events detected on the nodes contributed to the performance event.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

1. Display the **Event details** page to view information about the event.

2. Read the **Description**, which describes the workloads involved in the event and the cluster component in contention.

There is one victim volume whose latency was impacted by the cluster component in contention. The data processing node, which took over all workloads from its partner node, is the cluster component in contention. Under Component in Contention, the Data Processing icon is highlighted red and the name of the node that was handling data processing at the time of the event is displayed in parentheses.

3. In the **Description**, click the name of the victim volume.

The Performance/Volume Details page is displayed. At the bottom of the page, in the Events time line, a change event icon () indicates the time that Unified Manager detected the start of the HA takeover.

4. Point your cursor to the change event icon for the HA takeover.

Details about the HA takeover are displayed in the Events List table. In the Latency chart, an event indicates that the selected volume crossed the performance threshold due to high latency around the same time as the HA takeover.

- 5. Select Break down data by.
- 6. Under Latency, select Cluster Components.
- Click Submit.

The Cluster Components chart is displayed. The chart breaks down the total latency by cluster component.

- 8. At the bottom of the page, point your mouse cursor to the change event icon for the start of the HA takeover.
- 9. In the **Cluster Components** chart, compare the latency for data processing to the total latency in the **Latency** chart.

At the time of the HA takeover, there was a spike in data processing from the increased workload demand on the data processing node. The increased CPU utilization drove up the latency and triggered the event.

- 10. After fixing the failed node, use OnCommand System Manager to perform an HA giveback, which moves the workloads from the partner node to the fixed node.
- 11. After the HA giveback is complete, in Unified Manager, search for the event ID you recorded in Step 2.

The event triggered by the HA takeover is displayed on the Event details page. The event now has a state of obsolete, which indicates that the event is resolved.

12. In the **Description**, click the name of the victim volume.

The Performance/Volume Details page is displayed. At the bottom of the page, in the Events time line, a change event icon indicates the time that Unified Manager detected the completion of the HA giveback.

- 13. Select Break down data by.
- 14. Under Latency, select Cluster Components.

The Cluster Components chart is displayed.

15. At the bottom of the page, point your cursor to the change event icon for the HA giveback.

The change event is highlighted in the Events List table and indicates that the HA giveback was completed

successfully.

16. In the **Cluster Components** chart, compare the latency for data processing to the total latency in the **Latency** chart.

The latency at the data processing component has decreased, which has decreased the total latency. The node that the selected volume is now using for data processing has resolved the event.

Setting up a connection between a Unified Manager server and an external data provider

A connection between a Unified Manager server and an external data provider enables you to send cluster performance data to an external server so that storage managers can chart the performance metrics using third-party software.

A connection between a Unified Manager server and an external data provider is established through the menu option labeled "External Data Provider" in the maintenance console.

Performance data that can be sent to an external server

Unified Manager collects a variety of performance data from all the clusters that it is monitoring. You can send specific groups of data to an external server.

Depending on the performance data that you want to chart, you can choose to send one of the following groups of statistics:

Statistics group	Data included	Details
Performance Monitor	High-level performance statistics for the following objects: • LUNs • Volumes	This group provides total IOPS or latency for all LUNs and volumes in all monitored clusters. This group provides the smallest number of statistics.
Resource Utilization	Resource utilization statistics for the following objects: • Nodes • Aggregates	This group provides utilization statistics for the node and aggregate physical resources in all monitored clusters. It also provides the statistics collected in the Performance Monitor group.

Statistics group	Data included	Details
Drill Down	protocol statistics for all tracked per-protocol breakdowns to objects: per-protocol breakdowns to seven tracked object types monitored clusters.	This group provides read/write and per-protocol breakdowns for all seven tracked object types in all monitored clusters.
	NodesAggregatesLUNsVolumes	It also provides the statistics collected in the Performance Monitor group and in the Resource Utilization group.
	DisksLIFs	This group provides the largest number of statistics.
	Ports/NICs	



If the name of a cluster, or cluster object, is changed on the storage system, both the old and the new objects will contain performance data on the external server (called the "metric_path"). The two objects are not correlated as the same object. For example, if you change the name of a volume from "volume1_acct" to "acct_vol1", you will see old performance data for the old volume, and new performance data for the new volume.

See the Knowledge Base article for the list of all performance counters that can be sent to an external data provider.

Unified Manager performance counters that can be exported to an External Data Provider

Setting up Graphite to receive performance data from Unified Manager

Graphite is an open software tool for gathering and graphing performance data from computer systems. Your Graphite server and software must be configured correctly to receive statistical data from Unified Manager.

After you have installed Graphite according to the installation instructions, you need to make the following changes to support statistical data transfer from Unified Manager:

• In the /opt/graphite/conf/carbon.conf file, the maximum number of files that can be created on the Graphite server per minute must be set to 200 (MAX_CREATES_PER_MINUTE = 200).

Depending on the number of clusters in your configuration and the statistics objects you have selected to send, there might be thousands of new files that need to be created initially. At 200 files per minute it might take 15 minutes or longer before all metric files are initially created. After all the unique metric files have been created, this parameter is no longer relevant.

- If you are running Graphite on a server deployed using an IPv6 address, the value for LINE_RECEIVER_INTERFACE in the /opt/graphite/conf/carbon.conf file must be changed from "0.0.0.0" to "::" (LINE RECEIVER INTERFACE = ::)
- In the /opt/graphite/conf/storage-schemas.conf file, the retentions parameter must be used to set the frequency to 5 minutes and the retention period to the number of days that is relevant for your environment.

The retention period can be as long as what your environment allows, but the frequency value must be set to 5 minutes for at least one retention setting. In the following example, a section is defined for Unified Manager using the pattern parameter, and the values set the initial frequency to 5 minutes and the retention period to 100 days: pattern = ^netapp-performance\..*retentions = 5m:100d



If the default vendor tag is changed from "netapp-performance" to something different, that change must be reflected in the pattern parameter as well.



If the Graphite server is unavailable when the Unified Manager server is attempting to send performance data, the data is not sent and there will be a gap in collected data.

Configuring a connection from a Unified Manager server to an external data provider

Unified Manager can send cluster performance data to an external server. You can specify the type of statistical data that is sent, and the interval at which data is sent.

Before you begin

- You must have a user ID authorized to log in to the maintenance console of the Unified Manager server.
- · You must have the following information about the external data provider:
 - Server name or IP address (IPv4 or IPv6)
 - Server default port (if not using default port 2003)
- You must have configured the remote server and third-party software so that it can receive statistical data from the Unified Manager server.
- You must know which group of statistics you want to send:
 - PERFORMANCE INDICATOR: Performance monitor statistics
 - RESOURCE UTILIZATION: Resource utilization and Performance monitor statistics
 - DRILL DOWN: All statistics
- You must know the time interval at which you want to transmit statistics: 5, 10, or 15 minutes

By default, Unified Manager collects statistics at 5-minute intervals. If you set the transmit interval to 10 (or 15) minutes, the amount of data that is sent during each transmission is two (or three) times larger than when using the default 5-minute interval.



If you change the Unified Manager performance collection interval to 10 or 15 minutes, you must change the transmit interval so that it is equal to, or larger, than the Unified Manager collection interval.

About this task

You can configure a connection between one Unified Manager server and one external data provider server.

Steps

1. Log in as the maintenance user to the maintenance console of the Unified Manager server.

The Unified Managermaintenance console prompts are displayed.

2. In the maintenance console, type the number of the External Data Provider menu option.

The External Server Connection menu is displayed.

3. Type the number of the Add/Modify Server Connection menu option.

The current server connection information is displayed.

- 4. When prompted, type y to continue.
- 5. When prompted, enter the IP address or name of the destination server and the server port information (if different from the default port 2003).
- 6. When prompted, type y to verify that the information you entered is correct.
- 7. Press any key to return to the External Server Connection menu.
- 8. Type the number of the **Modify Server Configuration** menu option.

The current server configuration information is displayed.

- 9. When prompted, type y to continue.
- 10. When prompted, enter the type of statistics to send, the time interval at which the statistics are sent, and whether you want to enable the transmission of statistics now:

For	Enter
Statistics group ID	0 - PERFORMANCE_INDICATOR (default)1 - RESOURCE_UTILIZATION2 - DRILL_DOWN
Vendor tag	A descriptive name for the folder where the statistics will be stored on the external server. "netapp-performance" is the default name, but you can enter another value. By using dotted notation you can define a hierarchical folder structure. For example, by entering stats.performance.netapp the statistics will be located in stats > performance > netapp.
Transmit interval	5 (default), 10, or 15 minutes
Enable/disable	0 - Disable 1 - Enable (default)

11. When prompted, type y to verify that the information you entered is correct.

- 12. Press any key to return to the External Server Connection menu.
- 13. Type x to exit the maintenance console.

Results

After you have configured the connection, the selected performance data is sent to the destination server at the time interval you specified. It takes a few minutes before the metrics start to appear in Graphite. You might need to refresh your browser to see the new metrics in the metric hierarchy.

Monitor and manage cluster health

Introduction to OnCommand Unified Manager health monitoring

Unified Manager helps you to monitor a large number of systems running ONTAP software through a centralized user interface. The Unified Manager server infrastructure delivers scalability, supportability, and enhanced monitoring and notification capabilities.

The key capabilities of Unified Manager include monitoring, alerting, managing availability and capacity of clusters, managing protection capabilities, monitoring performance, configuring and managing of Infinite Volumes, annotating storage objects, and bundling of diagnostic data and sending it to technical support.

You can use Unified Manager to monitor your clusters. When issues occur in the cluster, Unified Manager notifies you about the details of such issues through events. Some events also provide you with a remedial action that you can take to rectify the issues. You can configure alerts for events so that when issues occur, you are notified through email, and SNMP traps.

You can use Unified Manager to manage storage objects in your environment by associating them with annotations. You can create custom annotations and dynamically associate clusters, storage virtual machines (SVMs), and volumes with the annotations through rules.

You can also plan the storage requirements of your cluster objects using the information provided in the capacity and health charts, for the respective cluster object.

Unified Manager health monitoring features

Unified Manager is built on a server infrastructure that delivers scalability, supportability, and enhanced monitoring and notification capabilities. Unified Manager supports monitoring of systems running ONTAP software.

Unified Manager includes the following features:

- Discovery, monitoring, and notifications for systems that are installed with ONTAP software:
 - Physical objects: nodes, disks, disk shelves, SFO pairs, ports, and Flash Cache
 - Logical objects: clusters, storage virtual machines (SVMs), aggregates, volumes, LUNs, namespaces, qtrees, LIFs, Snapshot copies, junction paths, NFS exports, CIFS shares, user and group quotas, and initiator groups
 - Protocols: CIFS, NFS, FC, iSCSI, NVMe, and FCoE
 - Storage efficiency: SSD aggregates, Flash Pool aggregates, FabricPool aggregates, deduplication, and compression
 - Protection: SnapMirror relationships (synchronous and asynchronous) and SnapVault relationships
- Viewing the cluster discovery and monitoring status
- MetroCluster configuration: viewing and monitoring the configuration, MetroCluster switches and bridges, issues, and connectivity status of the cluster components
- Enhanced alerts, events, and threshold infrastructure
- · LDAP, LDAPS, SAML authentication, and local user support

- RBAC (for a predefined set of roles)
- · AutoSupport and support bundle
- · Enhanced dashboard to show capacity, availability, protection, and performance health of the environment
- Volume move interoperability, volume move history, and junction path change history
- Scope of Impact area that graphically displays the resources that are impacted for events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events
- Possible Effect area that displays the effect of the MetroCluster events
- Suggested Corrective Actions area that displays the actions that can be performed to address events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events
- Resources that Might be Impacted area that displays the resources that might be impacted for events such as for the Volume Offline event, the Volume Restricted event, and the Thin-Provisioned Volume Space At Risk event
- · Support for SVMs with:
 - FlexVol volumes
 - FlexGroup volumes
 - Infinite Volumes
- · Support for monitoring node root volumes
- Enhanced Snapshot copy monitoring, including computing reclaimable space and deleting Snapshot copies
- · Annotations for storage objects
- Report creation and management of storage object information such as physical and logical capacity, utilization, space savings, and related events
- Integration with OnCommand Workflow Automation to execute workflows

The Storage Automation Store contains NetApp-certified automated storage workflow packs developed for use with OnCommand Workflow Automation (WFA). You can download the packs, and then import them to WFA to execute them. The automated workflows are available at the following Storage Automation Store

Unified Manager interfaces used to manage storage system health

This section has information about the two user interfaces that OnCommand Unified Manager provides for troubleshooting data storage capacity, availability, and protection issues. The two UIs are the Unified Manager web UI and the maintenance console.

If you want to use the protection features in Unified Manager, you must also install and configure OnCommand Workflow Automation (WFA).

Unified Manager web UI

The Unified Manager web UI enables an administrator to monitor and troubleshoot cluster issues relating to data storage capacity, availability, and protection.

This section describes some common workflows that an administrator can follow to troubleshoot storage capacity, data availability, or protection issues displayed in the Unified Manager web UI.

Maintenance console

The maintenance console enables an administrator to monitor, diagnose, and address operating system issues, version upgrade issues, user access issues, and network issues related to the Unified Manager server itself. If the Unified Manager web UI is unavailable, the maintenance console is the only form of access to Unified Manager.

This section provides directions for accessing the maintenance console and using it to resolve issues related to the functioning of the Unified Manager server.

Common Unified Manager health workflows and tasks

Some common administrative workflows and tasks associated with Unified Manager include selecting the storage clusters that are to be monitored; diagnosing conditions that adversely affect data availability, capacity, and protection; creating protection relationships; restoring lost data; configuring and managing Infinite Volumes; and bundling and sending diagnostic data to technical support (when necessary).

Unified Manager enables storage administrators to view a dashboard, assess the overall capacity, availability, and protection health of the managed storage clusters, and then quickly identify, locate, diagnose, and assign for resolution any specific issues that might arise.

The most important issues related to a cluster, storage virtual machine (SVM), volume, FlexGroup volume, or protection relationship that affect the storage capacity, data availability, or protection reliability of your managed storage objects are displayed in the system health graphs and events on the Dashboards/Overview page. When critical issues are identified, the this page provides links to support appropriate troubleshooting workflows.

Unified Manager can also be included in workflows that include related manageability tools—such as OnCommand Workflow Automation (WFA)--to support the direct configuration of storage resources.

Common workflows related to the following administrative tasks are described in this document:

· Diagnosing and managing availability issues

If hardware failure or storage resource configuration issues cause the display of data availability events in the Dashboards/Overview page, storage administrators can follow the embedded links to view connectivity information about the affected storage resource, view troubleshooting advice, and assign issue resolution to other administrators.

· Configuring and monitoring performance incidents

The OnCommand Administrator can monitor and manage the performance of the storage system resources that are being monitored. See the *Unified Manager Workflow Guide for Managing Cluster Performance* for more information.

· Diagnosing and managing volume capacity issues

If volume storage capacity issues are displayed in the Dashboards/Overview page, storage administrators can follow the embedded links to view the current and historical trends related to the storage capacity of the affected volume, view troubleshooting advice, and assign issue resolution to other administrators.

· Configuring, monitoring, and diagnosing protection relationship issues

After creating and configuring protection relationships, storage administrators can view the potential issues related to protection relationships in the Dashboards/Overview page, and they can follow the embedded links to view the current state of the protection relationships, the current and historical protection job success information about the affected relationships, and troubleshooting advice, and to assign issue resolution to other administrators. Storage administrators can also configure and manage SnapMirror and SnapVault relationships.

- · Creating backup files and restoring data from backup files.
- · Associating storage objects with annotations

By associating storage objects with annotations, storage administrators can filter and view the events that are related to the storage objects, which enables storage administrators to prioritize and resolve the issues that are associated with the events.

Sending a support bundle to technical support

Storage administrators can retrieve and send a support bundle to technical support by using the maintenance console. Support bundles must be sent to technical support when the issue requires more detailed diagnosis and troubleshooting than what an AutoSupport message provides.

· Creating new reports for import

Storage administrators can create new .rptdesign files by using the Eclipse plug-in for Business Intelligence and Reporting Tools (BIRT). These reports can be imported to the Unified Manager UI and viewed in the Reports page.

The reports that are displayed on the Reports page provide the current status of the storage objects. You can make important decisions—such as decisions about storage procurement—based on the current usage. These reports provide a detailed view of storage objects such as volumes, disk shelves, and aggregates.

The Reports page in the Unified Manager UI enables you to view detailed information about the reports that you generate. You can search for a specific report, save a report, and delete a report from the Reports page. You can also schedule, share, and import a report from this page.

· Creating, configuring, monitoring, and protecting Infinite Volumes

After using the Workflow Automation tool to create, configure, and define storage classes for an Infinite Volume, storage administrators can use Unified Manager to monitor, set notification thresholds, and define the data policy for that volume and its storage classes. Optionally, storage administrators can use WFA and Unified Manager to set up data protection for the Infinite Volume.

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.

Resolving a flash card offline condition

This workflow provides an example of how you might resolve a flash card offline condition. In this scenario, you are an administrator or operator monitoring the dashboard to check for problems with availability. You see a flash card offline condition and you want to determine the possible cause of and resolution to the problem.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

The event information and links displayed in the Availability area of the Unified ManagerDashboards/Overview page monitor the overall availability of data storage resources on the monitored clusters enable you to diagnose specific events that might affect that availability.

In this scenario, the Dashboards/Overview page displays the event Flash Cards Offline in its Availability Incidents section. If a flash card is offline, availability of stored data is impeded because the performance of the cluster node on which it is installed is impaired. You can perform the following steps to localize and identify the potential problem:

Steps

 From the Availability panel in the Unresolved Incidents and Risks section, click the hypertext link displayed for Flash Cards Offline.

The Event details page for the availability incident is displayed.

- On the Event details page, you can review the information displayed in the Cause field and perform one or more of the following tasks:
 - Assign the event to an administrator. Assigning events
 - Click the source of the event, in this case the cluster node on which the offline flash card is located, to get more information about that node. Performing corrective action for a flash card offline
 - Acknowledge the event. Acknowledging and resolving events

Performing corrective action for a flash card offline

After reviewing the description in the Cause field of the Flash Card Offline Event details page, you can search for additional information helpful to resolving the condition.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

In this example scenario, the event summary provided on the Event details page contains the following information about the offline flash card condition:

```
Severity: Critical
State: New
Impact Level: Incident
Impact Area: Availability
Source: alpha-node
Source Type: Node
Acknowledged By:
Resolved By:
Assigned To:
Cause: Flash cards at slot numbers 3 are offline.
Alert Settings:
```

The event information indicates that the flash card installed in slot 3 in the cluster node named "alpha-node" is offline.

The information localizes the flash card offline condition to a specific slot on a specific cluster node but does not suggest a reason that the flash card is offline.

Steps

1. To obtain further details that might help you diagnose the flash card offline condition, you can click the name of the source of the event.

In this example, the source of the event is the "alpha-node" cluster node. Clicking that node name displays the HA Details on the Nodes tab of the Health/Cluster details page for the affected cluster. The displayed HA Details displays information about the HA pair to which that node belongs.

In this example, the relevant information is in the Events summary table on the HA Details. The table specifies the flash card offline event, the time the event was generated, and, again, the cluster node from which this event originated.

2. Using the ONTAP CLI or OnCommand System Manager, access the Event Manager System (EMS) logs for the affected cluster.

In this example, you use the event name, the event time, and the event source to find the EMS report on this event. The EMS report on the event contains a detailed description of the event and often advice to remedy the condition indicated by the event.

After you finish

After you diagnose the problem, contact the appropriate administrator or operator to complete the manual steps necessary to get the flash card back online.

Scanning for and resolving storage failover interconnect link down conditions

This workflow provides an example of how you might scan for, evaluate, and resolve downed storage failover interconnect link conditions. In this scenario, you are an administrator using Unified Manager to scan for storage failover risks before starting an ONTAP version upgrade on your nodes.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

If storage failover interconnections between HA pair nodes fail during a nondisruptive upgrade attempt, the upgrade fails. Therefore, common practice is for the administrator to monitor and confirm storage failover reliability on the cluster nodes targeted for upgrade before the start of an upgrade.

Steps

- 1. To check for recent availability events related to storage failover issues, check the Availability Incidents section and the Availability Risks listings on the **Dashboards/Overview** page.
- 2. To check further for all availability events related to storage failover issues, perform the following steps:
 - a. Click the Availability Incidents link on the Dashboards/Overview page.

The Events inventory page displays all events on the monitored clusters.

- b. On the **Events** inventory page, select the options **Incident** and **Risk** in the Filter column.
- c. At the top of the **Events** inventory page Names column, click = and enter *failover in the text box to limit the event to display to storage failover-related events.

All past events related to storage failover conditions are displayed.

In this scenario, the Unified Manager displays the event, "Storage Failover Interconnect One or More Links Down" in its Availability Incidents section.

- 3. If one or more events related to storage failover are displayed either on the **Dashboards/Overview** page or on the **Events** inventory page, perform the following steps:
 - a. Click the event title link to display event details for that event.

In this example, you click the event title "Storage Failover Interconnect One or More Links Down".

The Event details page for that event is displayed.

- b. On the **Event** details page, you can perform one or more of the following tasks:
 - Review the error message in the Cause field and evaluate the issue. Performing corrective action for storage failover interconnect links down
 - Assign the event to an administrator. Assigning events
 - Acknowledge the event. Acknowledging and resolving events

Performing corrective action for storage failover interconnect links down

When you display the Event details page of a storage failover-related event, you can review the summary information of the page to determine the urgency of the event, possible cause of the issue, and possible resolution to the issue.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

In this example scenario, the event summary provided on the Event details page contains the following information about the storage failover interconnect link down condition:

```
Event: Storage Failover Interconnect One or More Links Down

Summary

Severity: Warning
State: New
Impact Level: Risk
Impact Area: Availability
Source: aardvark
Source Type: Node
Acknowledged By:
Resolved By:
Assigned To:
Cause: At least one storage failover interconnected link
between the nodes aardvark and bonobo is down.
RDMA interconnect is up (Link0 up, Link1 down)
```

The example event information indicates that a storage failover interconnect link, Link1, between HA pair nodes aardvark and bonobo is down, but that Link0 between Apple and Boy is active. Because one link is active, the remote dynamic memory access (RDMA) is still functioning and a storage failover job can still succeed.

However, to ensure against both links failing and storage failover protection being totally disabled, you decide to further diagnose the reason for Link1 going down.

Steps

1. From the **Event** details page, you can click the link to the event specified in the Source field to obtain further details of other events that might be related to the storage failover interconnection link down condition.

In this example, the source of the event is the node named aardvark. Clicking that node name displays the HA Details for the affected HA pair, aardvark and bonobo, on the Nodes tab of the Health/Cluster details page, and displays other events that recently occurred on the affected HA pair.

2. Review the **HA Details** for more information relating to the event.

In this example, the relevant information is in the Events table. The table shows the "Storage Failover Connection One or More Link Down" event, the time the event was generated, and, again, the node from which this event originated.

After you finish

Using the node location information in the HA Details, request or personally complete a physical inspection and repair of the storage failover issue on the affected HA pair nodes.

Resolving volume offline issues

This workflow provides an example of how you might evaluate and resolve a volume offline event that Unified Manager might display in the Availability area of the Dashboards/Overview page. In this scenario, you are an administrator using Unified Manager to troubleshoot one or more volume offline events that are displayed on the Dashboards/Overview page.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

Volumes might be reported offline for several reasons:

- The SVM administrator has deliberately taken the volume offline.
- The volume's hosting cluster node is down and storage failover to its HA pair partner has failed also.
- The volume's hosting storage virtual machine (SVM) is stopped because the node hosting the root volume of that SVM is down.
- The volume's hosting aggregate is down due to simultaneous failure of two RAID disks.

You can use the Dashboards/Overview page and the Health/Cluster, Health/SVM, and Health/Volume details pages to confirm or eliminate one or more of these possibilities.

Steps

1. From the **Availability** panel in the **Unresolved Incidents and Risks** section, click the hypertext link displayed for the Volume Offline event.

The Event details page for the availability incident is displayed.

- 2. On that page, check the notes for any indication that the SVM administrator has taken the volume in question offline.
- On the Event details page, you can review the information for one or more of the following tasks:
 - Review the information displayed in the Cause field for possible diagnostic guidance.

In this example, the information in the Cause field informs you only that the volume is offline.

- Check the Notes and Updates area for any indication that the SVM administrator has deliberately taken the volume in question offline.
- Click the source of the event, in this case the volume that is reported offline, to get more information about that volume. Performing corrective action for volume offline conditions
- Assign the event to an administrator. Assigning events
- · Acknowledge the event or, if appropriate, mark it as resolved. Acknowledging and resolving events

Performing diagnostic actions for volume offline conditions

After navigating to the Health/Volume details page of a volume reported to be offline, you can search for additional information helpful to diagnosing the volume offline condition.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

If the volume that is reported offline was not taken offline deliberately, that volume might be offline for several reasons.

Starting at the offline volume's Health/Volume details page, you can navigate to other pages and panes to confirm or eliminate possible causes:

Choices

• Click **Health/Volume** details page links to determine if the volume is offline because its host node is down and storage failover to its HA pair partner has failed also.

See Determining if a volume offline condition is caused by a down node.

• Click **Health/Volume** details page links to determine if the volume is offline and its host storage virtual machine (SVM) is stopped because the node hosting the root volume of that SVM is down.

See Determining if a volume is offline and SVM is stopped because a node is down.

• Click **Health/Volume** details page links to determine if the volume is offline because of broken disks in its host aggregate.

See Determining if a volume is offline because of broken disks in an aggregate.

Determining if a volume is offline because its host node is down

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because its host node is down and that storage failover to its HA pair partner is unsuccessful.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

To determine if the volume offline condition is caused by failure of the hosting node and subsequent unsuccessful storage failover, perform the following actions:

Steps

1. Locate and click the hypertext link displayed under SVM in the **Related Devices** pane of the offline volume's **Health/Volume** details page.

The Health/Storage Virtual Machine details page displays information about the offline volume's hosting storage virtual machine (SVM).

2. In the **Related Devices** pane of the **Health/Storage Virtual Machine** details page, locate and click hypertext link displayed under Volumes.

The Health/Volumes inventory page displays a table of information about all the volumes hosted by the SVM.

3. On the **Health/Volumes** inventory page State column header, click the filter symbol =, and then select the option **Offline**.

Only the SVM volumes that are in offline state are listed.

4. On the **Health/Volumes** inventory page, click the grid symbol **,** and then select the option **Cluster Nodes**.

You might need to scroll in the grid selection box to locate the Cluster Nodes option.

The Cluster Nodes column is added to the volumes inventory and displays the name of the node that hosts each offline volume.

5. On the **Health/Volumes** inventory page, locate the listing for the offline volume and, in its Cluster Node column, click the name of its hosting node.

The Nodes tab on the Health/Cluster details page displays the state of the HA pair of nodes to which the hosting node belongs. The state of the hosting node and the success of any cluster failover operation is indicated in the display.

After you finish

After you confirm that the volume offline condition exists because its host node is down and storage failover to the HA pair partner has failed, contact the appropriate administrator or operator to manually restart the down node and fix the storage failover problem.

Determining if a volume is offline and its SVM is stopped because a node is down

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because its host storage virtual machine (SVM) is stopped due to the node hosting the root volume of that SVM being down.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

To determine if the volume offline condition is caused its host SVM being stopped because the node hosting the root volume of that SVM is down, perform the following actions:

Steps

- 1. Locate and click the hypertext link displayed under SVM in the **Related Devices** pane of the offline volume's **Health/Volume** details page.
- 2. Locate and click the hypertext link displayed under the SVM in the **Related Devices** pane of the offline volume's **Health/Volume** details page.

The Health/Storage Virtual Machine details page displays the "running" or the "stopped" status of the hosting SVM. If the SVM status is running, then the volume offline condition is not caused by the node

hosting the root volume of that SVM being down.

- If the SVM status is stopped, then click View SVMs to further identify the cause of the hosting SVM being stopped.
- 4. On the **Health/Storage Virtual Machines** inventory pageSVM column header, click the filter symbol = and then type the name of the stopped SVM.

The information for that SVM is shown in a table.

5. On the **Health/Storage Virtual Machines** inventory page, click and then select the option **Root Volume**.

The Root Volume column is added to the SVM inventory and displays the name of the root volume of the stopped SVM.

6. In the Root Volume column, click the name of the root volume to display the **Health/Storage Virtual Machine** details page for that volume.

If the status of the SVM root volume is (Online), then the original volume offline condition is not caused because the node hosting the root volume of that SVM is down.

- 7. If the status of the SVM root volume is (Offline), then locate and click the hypertext link displayed under Aggregate in the **Related Devices** pane of the SVM root volume's **Health/Volume** details page.
- 8. Locate and click the hypertext link displayed under Node in the **Related Devices** pane of the Aggregate's **Health/Aggregate** details page.

The Nodes tab on the Health/Cluster details page displays the state of the HA pair of nodes to which the SVM root volume's hosting node belongs. The state of the node is indicated in the display.

After you finish

After you confirm that the volume offline condition is caused by that volume's host SVM offline condition, which itself is caused by the node that hosts the root volume of that SVM being down, contact the appropriate administrator or operator to manually restart the down node.

Determining if a volume is offline because of broken disks in an aggregate

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because RAID disk problems have taken its host aggregate offline.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

To determine if the volume offline condition is caused by RAID disk problems that are taking the hosting aggregate offline, perform the following actions:

Steps

1. Locate and click the hypertext link displayed under Aggregate in the **Related Devices** pane of the **Health/Volume** details page.

The Health/Aggregate details page displays the online or offline status of the hosting aggregate. If the aggregate status is online, then RAID disk problems are not the cause of the volume being offline.

- 2. If the aggregate status is offline, then click **Disk Information** and look for broken disk events in the **Events** list on the **Disk Information** tab.
- 3. To further identify the broken disks, click the hypertext link displayed under Cluster in the **Related Devices** pane.

The Health/Cluster details page is displayed.

4. Click **Disks**, and then select **Broken** in the **Filters** pane to list all disks in the broken state.

If the disks in the broken state caused the offline state of the host aggregate, the name of the aggregate is displayed in the Impacted Aggregate column.

After you finish

After confirming that the volume offline condition is caused by broken RAID disks and the consequent offline host aggregate, contact the appropriate administrator or operator to manually replace the broken disks and put the aggregate back online.

Resolving capacity issues

This workflow provides an example of how you can resolve a capacity issue. In this scenario, you are an administrator or operator and you access the Unified ManagerDashboards/Overview page to see if any of the monitored storage objects have capacity issues. You see that there is a volume with a capacity risk, and you want to determine the possible cause of and resolution to the problem.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role..

About this task

On the Dashboards/Overview page, you look at the Unresolved Incidents and Risks area and see a "Volume Space Full" error event in the Capacity pane under SVM Volume Capacity at Risk.

Steps

1. In the **Unresolved Incidents and Risks** area of the **Dashboards/Overview** page, click the name of the Volume Space Full error event in the **Capacity** pane.

The Event details page for the error is displayed.

- 2. From the **Event** details page, you can perform one or more of the following tasks:
 - Review the error message in the Cause field and click the suggestions under Suggested Remedial
 Actions to review descriptions of possible remediations. Performing suggested remedial actions for a
 full volume
 - Click the object name, in this case a volume, in the Source field to get details about the object. Volume details page

- Look for notes that might have been added about this event. Adding and reviewing notes associated with an event
- · Add a note to the event. Adding and reviewing notes associated with an event
- Assign the event to another user. Assigning events
- Acknowledge the event. Acknowledging and resolving events
- Mark the event as resolved. Acknowledging and resolving events

Performing suggested remedial actions for a full volume

After receiving a "Volume Space Full" error event, you review the suggested remedial actions on the Event details page and decide to perform one of the suggested actions.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

A user with any role can perform all of the tasks in this workflow that use Unified Manager.

About this task

In this example, you have seen a Volume Space Full error event on the Unified ManagerDashboards/Overview page and have clicked the name of the event.

Possible remedial actions you might perform for a full volume include the following:

- · Enabling autogrow, deduplication, or compression on the volume
- · Resizing or moving the volume
- · Deleting or moving data from the volume

Although all of these actions must be performed from either OnCommand System Manager or the ONTAP CLI, you can use Unified Manager to find information you might need to determine which actions to take.

Steps

- 1. From the **Event** details page, you click the volume name in the Source field to view details about the affected volume.
- On the Health/Volume details page, you click Configuration and see that deduplication and compression are already enabled on the volume.

You decide to resize the volume.

- 3. In the **Related Devices** pane, you click the name of the hosting aggregate to see if the aggregate can accommodate a larger volume.
- 4. On the **Health/Aggregate** details page, you see that the aggregate hosting the full volume does have enough uncommitted capacity, so you use OnCommand System Manager to resize the volume, giving it more capacity.

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.

Types of SnapMirror protection

Depending on the deployment of your data storage topology, Unified Manager enables you to configure multiple types of SnapMirror protection relationships. All variations of SnapMirror protection offer failover disaster recovery protection, but offer differing capabilities in performance, version flexibility, and multiple backup copy protection.

Traditional SnapMirror Asynchronous protection relationships

Traditional SnapMirror Asynchronous protection provides block replication mirror protection between source and destination volumes.

In traditional SnapMirror relationships, mirror operations execute faster than they would in alternative SnapMirror relationships because the mirror operation is based on block replication. However, traditional SnapMirror protection requires that the destination volume run under the same or later minor version of ONTAP software as the source volume within the same major release (for example, version 8.x to 8.x, or 9.x to 9.x).

SnapMirror Asynchronous protection with version-flexible replication

SnapMirror Asynchronous protection with version-flexible replication provides logical replication mirror protection between source and destination volumes, even if those volumes are running under different versions of ONTAP 8.3 or later software (for example, version 8.3 to 8.3, or 8.3 to 9.1, or 9.0 to 8.3).

In SnapMirror relationships with version-flexible replication, mirror operations do not execute as quickly as they would in traditional SnapMirror relationships.

Because of slower execution, SnapMirror with version-flexible replication protection is not suitable to implement in either of the following circumstances:

- The source object contains more than 10 million files to protect.
- The recovery point objective for the protected data is two hours or less. (That is, the destination must always contain mirrored, recoverable data that is no more than two hours older than data at the source.)

In either of the listed circumstances, the faster block-replication based execution of default SnapMirror protection is required.

SnapMirror Asynchronous protection with version-flexible replication and backup option

SnapMirror Asynchronous protection with version-flexible replication and backup option provides mirror protection between source and destination volumes and the capability to store multiple copies of the mirrored data at the destination.

The storage administrator can specify which Snapshot copies are mirrored from source to destination and can also specify how long to retain those copies at the destination, even if they are deleted at the source.

In SnapMirror relationships with version-flexible replication and backup option, mirror operations do not execute as quickly as they would in traditional SnapMirror relationships.

SnapMirror Synchronous protection with strict synchronization

SnapMirror Synchronous protection with "strict" synchronization ensures that the primary and secondary volumes are always a true copy of each other. If a replication failure occurs when attempting to write data to the secondary volume, then the client I/O to the primary volume is disrupted.

SnapMirror Synchronous protection with regular synchronization

SnapMirror Synchronous protection with "regular" synchronization does not require that the primary and secondary volume are always a true copy of each other; thereby ensuring availability of the primary volume. If a replication failure occurs when attempting to write data to the secondary volume, the primary and secondary volumes fall out of sync and client I/O will continue to the primary volume.



The Restore button and the Relationship operation buttons are not available when monitoring synchronous protection relationships from the Health/Volumes inventory page or the Health/Volume details page.

Setting up protection relationships in Unified Manager

There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have established peer relationships between two clusters or two storage virtual machines (SVMs).
- OnCommand Workflow Automation must be integrated with Unified Manager:
 - Set up OnCommand Workflow Automation
 - Verifying Unified Manager data source caching in Workflow Automation

Steps

- 1. Depending on the type of protection relationship you want to create, do one of the following:
 - Create a SnapMirror protection relationship.
 - Create a SnapVault protection relationship.
- 2. If you want to create a policy for the relationship, depending on the relationship type you are creating, do one of the following:
 - Create a SnapVault policy.
 - · Create a SnapMirror policy.
- 3. Create a SnapMirror or SnapVault schedule.

Configuring a connection between Workflow Automation and Unified Manager

You can configure a secure connection between OnCommand Workflow Automation (WFA) and Unified Manager. Connecting to Workflow Automation enables you to use protection features such as SnapMirror and SnapVault configuration workflows, as well as

commands for managing SnapMirror relationships.

Before you begin

- The installed version of Workflow Automation must be 4.2 or greater.
- You must have installed "WFA pack for managing Clustered Data ONTAP" version 9.5.0 or greater on the WFA server. You can download the required pack from the NetAppStorage Automation Store.

WFA pack for managing ONTAP

 You must have the name of the database user that you created in Unified Manager to support WFA and Unified Manager connections.

This database user must have been assigned the Integration Schema user role.

- You must be assigned either the Administrator role or the Architect role in Workflow Automation.
- You must have the host address, port number 443, user name, and password for the Workflow Automation setup.
- You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click **Workflow Automation** in the left Setup menu.
- In the OnCommand Unified Manager Database User area of the Setup/Workflow Automation page, select the name, and enter the password for the database user that you created to support Unified Manager and Workflow Automation connections.
- 3. In the **OnCommand Workflow Automation Credentials** area of the **Setup/Workflow Automation** page, enter the host name or IP address (IPv4 or IPv6), and the user name and password for the Workflow Automation setup.

You must use the Unified Manager server port (port 443).

- 4. Click Save.
- 5. If you use a self-signed certificate, click Yes to authorize the security certificate.

The Setup/Workflow Automation page displays.

6. Click **Yes** to reload the web UI, and add the Workflow Automation features.

Verifying Unified Manager data source caching in Workflow Automation

You can determine whether Unified Manager data source caching is working correctly by checking if data source acquisition is successful in Workflow Automation. You might do this when you integrate Workflow Automation with Unified Manager to ensure that Workflow Automation functionality is available after the integration.

Before you begin

You must be assigned either the Administrator role or the Architect role in Workflow Automation to perform this task.

Steps

- 1. From the Workflow Automation UI, select Execution > Data Sources.
- 2. Right-click the name of the Unified Manager data source, and then select Acquire Now.
- 3. Verify that the acquisition succeeds without errors.

Acquisition errors must be resolved for Workflow Automation integration with Unified Manager to succeed.

Creating a SnapMirror protection relationship from the Health/Volume details page

You can use the Health/Volume details page to create a SnapMirror relationship so that data replication is enabled for protection purposes. SnapMirror replication enables you to restore data from the destination volume in the event of data loss on the source.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- · If the volume is a FlexGroup volume
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, right-click in the topology view the name of a volume that you want to protect.
- 2. Select **Protect > SnapMirror** from the menu.

The Configure Protection dialog box is displayed.

- 3. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
- 4. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
- 5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
- 6. Click Apply.

You are returned to the Health/Volume details page.

7. Click the protection configuration job link at the top of the **Health/Volume** details page.

The job's tasks and details are displayed in the Protection/Job details page.

- 8. In the **Protection/Job** details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
- 9. When the job tasks are complete, click **Back** on your browser to return to the **Health/Volume** details page.

The new relationship is displayed in the Health/Volume details page topology view.

Results

Depending on the destination SVM you specified during configuration or on the options you enabled in your Advanced settings, the resulting SnapMirror relationship might be one of several possible variations:

- If you specified a destination SVM that runs under the same or a newer version of ONTAP compared to that of the source volume, a block-replication-based SnapMirror relationship is the default result.
- If you specified a destination SVM that runs under the same or a newer version of ONTAP (version 8.3 or higher) compared to that of the source volume, but you enabled version-flexible replication in the Advanced settings, a SnapMirror relationship with version-flexible replication is the result.
- If you specified a destination SVM that runs under an earlier version of ONTAP 8.3, or a version that is higher than that of the source volume and the earlier version supports version-flexible replication, a SnapMirror relationship with version-flexible replication is the automatic result.

Creating a SnapVault protection relationship from the Health/Volume details page

You can create a SnapVault relationship using the Health/Volume details page so that data backups are enabled for protection purposes on volumes.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up Workflow Automation to perform this task.

About this task

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, right-click a volume in the topology view that you want to protect.
- 2. Select **Protect** > **SnapVault** from the menu.

The Configure Protection dialog box is launched.

- 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.
- Click Apply.

You are returned to the Health/Volume details page.

7. Click the protection configuration job link at the top of the **Health/Volume** details page.

The Protection/Job details page is displayed.

8. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

When the job tasks are complete, the new relationships are displayed in the Health/Volume details page topology view.

Creating a SnapVault policy to maximize transfer efficiency

You can create a new SnapVault policy to set the priority for a SnapVault transfer. You use policies to maximize the efficiency of transfers from the primary to the secondary in a protection relationship.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.
- You must have already completed Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

The SnapVault tab is displayed.

- 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. 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.

Creating a SnapMirror policy to maximize transfer efficiency

You can create a SnapMirror policy to specify the SnapMirror transfer priority for protection relationships. SnapMirror policies enable you to maximize transfer efficiency from the source to the destination by assigning priorities so that lower-priority transfers are scheduled to run after normal-priority transfers.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.
- This task assumes that you have already completed the Destination Information area in the Configure Protection dialog box.

Steps

 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.

Creating SnapMirror and SnapVault schedules

You can create basic or advanced SnapMirror and SnapVault schedules to enable automatic data protection transfers on a source or primary volume so that transfers take place more frequently or less frequently, depending on how often the data changes on your volumes.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role...
- You must have already completed the Destination Information area in the Configure Protection dialog box.
- You must have set up Workflow Automation to perform this task.

Steps

1. From the **SnapMirror** tab or **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Schedule** link in the **Relationship Settings** area.

The Create Schedule dialog box is displayed.

- 2. In the **Schedule Name** field, type the name you want to give to the schedule.
- 3. Select one of the following:
 - Basic

Select if you want to create a basic interval-style schedule.

Advanced

Select if you want to create a cron-style schedule.

Click Create.

The new schedule is displayed in the SnapMirror Schedule or SnapVault Schedule drop-down list.

Performing a protection relationship failover and failback

When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation to perform this operation.

Steps

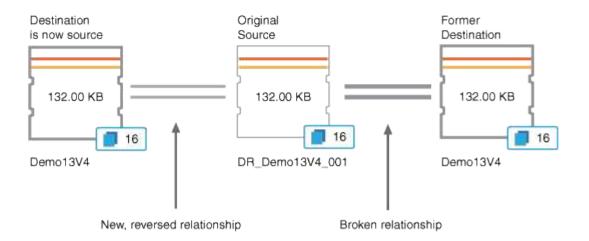
1. Break the SnapMirror relationship.

You must break the relationship before you can convert the destination from a data protection volume to a read/write volume, and before you can reverse the relationship.

2. Reverse the protection relationship.

When the original source volume is available again, you might decide to reestablish the original protection relationship by restoring the source volume. Before you can restore the source, you must synchronize it with the data written to the former destination. You use the reverse resync operation to create a new protection relationship by reversing the roles of the original relationship and synchronizing the source volume with the former destination. A new baseline Snapshot copy is created for the new relationship.

The reversed relationship looks similar to a cascaded relationship:



Break the reversed SnapMirror relationship.

When the original source volume is resynchronized and can again serve data, use the break operation to break the reversed relationship.

4. Remove the relationship.

When the reversed relationship is no longer required, you should remove that relationship before reestablishing the original relationship.

5. Resynchronize the relationship.

Use the resynchronize operation to synchronize data from the source to the destination and to reestablish the original relationship.

Breaking a SnapMirror relationship from the Health/Volume details page

You can break a protection relationship from the Health/Volume details page and stop data transfers between a source and destination volume in a SnapMirror relationship. You might break a relationship when you want to migrate data, for disaster recovery, or for application testing. The destination volume is changed to a read-write volume. You cannot break a SnapVault relationship.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, select from the topology the SnapMirror relationship you want to break.
- 2. Right-click the destination and select **Break** from the menu.

The Break Relationship dialog box is displayed.

- 3. Click **Continue** to break the relationship.
- 4. In the topology, verify that the relationship is broken.

Reversing protection relationships from the Health/Volume details page

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.
- The relationship must not be a SnapVault relationship.
- · A protection relationship must already exist.

- The protection relationship must be broken.
- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.

About this task

- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on the reverse resynchronization relationship are the same as those on the original protection relationship.

If policies and schedules do not exist, they are created.

Steps

- 1. From the **Protection** tab of the **Health/Volume** details page, locate in the topology the SnapMirror relationship on which you want to reverse the source and destination, and right-click it.
- 2. Select Reverse Resync from the menu.

The Reverse Resync dialog box is displayed.

3. Verify that the relationship displayed in the **Reverse Resync** dialog box is the one for which you want to perform the reverse resynchronization operation, and then click **Submit**.

The Reverse Resync dialog box is closed and a job link is displayed at the top of the Health/Volume details page.

4. Click **View Jobs** on the **Health/Volume** details page to track the status of each reverse resynchronization job.

A filtered list of jobs is displayed.

5. Click the Back arrow on your browser to return to the **Health/Volume** details page.

The reverse resynchronization operation is finished when all job tasks are completed successfully.

Removing a protection relationship from the Health/Volume details page

You can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- · You must have set up Workflow Automation.

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, select from the topology the SnapMirror relationship you want to remove.
- 2. Right-click the name of the destination and select **Remove** from the menu.

The Remove Relationship dialog box is displayed.

3. Click **Continue** to remove the relationship.

The relationship is removed from the Health/Volume details page.

Resynchronizing protection relationships from the Health/Volume details page

You can resynchronize data on a SnapMirror or SnapVault relationship that was broken and then the destination was made read/write so that data on the source matches the data on the destination. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation.

Steps

- 1. From the **Protection** tab of the **Health/Volume** details page, locate in the topology the protection relationship that you want to resynchronize and right-click it.
- Select Resynchronize from the menu.

Alternatively, from the **Actions** menu, select **Relationship** > **Resynchronize** to resynchronize the relationship for which you are currently viewing the details.

The Resynchronize dialog box is displayed.

- 3. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
- 4. Click Source Snapshot Copies; then, in the Snapshot Copy column, click Default.

The Select Source Snapshot Copy dialog box is displayed.

- 5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
- 6. Click Submit.

You are returned to the Resynchronize dialog box.

- 7. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
- 8. Click **Submit** to begin the resynchronization job.

The resynchronization job is started, you are returned to the Health/Volume details page and a jobs link is displayed at the top of the page.

9. Click View Jobs on the Health/Volume details page to track the status of each resynchronization job.

A filtered list of jobs is displayed.

10. Click the Back arrow on your browser to return to the **Health/Volume** details page.

The resynchronization job is finished when all job tasks successfully complete.

Resolving a protection job failure

This workflow provides an example of how you might identify and resolve a protection job failure from the Unified Manager dashboard.

Before you begin

Because some tasks in this workflow require that you log in using the OnCommand Administrator role, you must be familiar with the roles required to use various functionality, as described in Unified Manager user roles and capabilities.

About this task

In this scenario, you access the Dashboards/Overview page to see if there are any issues with your protection jobs. In the Protection Incident area, you notice that there is a Job Terminated incident, showing a Protection Job Failed error on a volume. You investigate this error to determine the possible cause and potential resolution.

Steps

1. In the **Protection Incidents** panel of the Dashboard **Unresolved Incidents and Risks** area, you click the **Protection job failed** event.



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.

Identifying the problem and performing corrective actions for a failed protection job

You review the job failure error message in the Cause field on the Event details page and determine that the job failed because of a Snapshot copy error. You then proceed to the Health/Volume details page page to gather more information.

Before you begin

You must have the OnCommand Administrator role.

About this task

The error message provided in the Cause field on the Event details page contains the following text about the failed job:

```
Protection Job Failed. Reason: (Transfer operation for relationship 'cluster2_src_svm:cluster2_src_vol2->cluster3_dst_svm: managed_svc2_vol3' ended unsuccessfully. Last error reported by Data ONTAP: Failed to create Snapshot copy 0426cluster2_src_vol2snap on volume cluster2_src_svm:cluster2_src_vol2. (CSM: An operation failed due to an ONC RPC failure.).)
*Job Details*
```

This message provides the following information:

· A backup or mirror job did not complete successfully.

The job involved a protection relationship between the source volume <code>cluster2_src_vol2</code> on the virtual server <code>cluster2_src_svm</code> and the destination volume <code>managed_svc2_vol3</code> on the virtual server <code>named cluster3_dst_svm</code>.

• A Snapshot copy job failed for 0426cluster2_src_vol2snap on the source volume cluster2 src svm:/cluster2 src vol2.

In this scenario, you can identify the cause and potential corrective actions of the job failure. However, resolving the failure requires that you access either the System Manager web UI or the ONTAP CLI commands.

Steps

1. You review the error message and determine that a Snapshot copy job failed on the source volume, indicating that there is probably a problem with your source volume.

Optionally, you could click the **Job Details** link at the end of the error message, but for the purposes of this scenario, you choose not to do that.

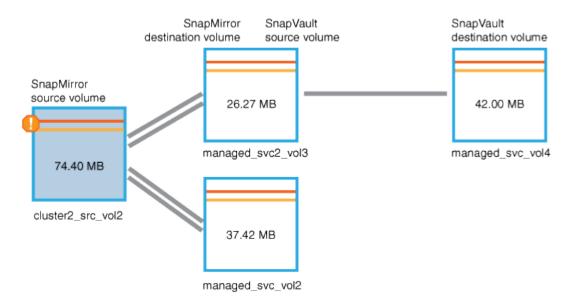
- 2. You decide that you want to try to resolve the event, so you do the following:
 - a. Click the **Assign To** button and select **Me** from the menu.
 - b. Click the **Acknowledge** button so that you do not continue to receive repeat alert notifications, if alerts were set for the event.
 - c. Optionally, you can also add notes about the event.
- Click the Source field in the Summary pane to see details about the source volume.

The **Source** field contains the name of the source object: in this case, the volume on which the Snapshot copy job was scheduled.

The Health/Volume details page displays for cluster2_src_vol2, showing the content of the Protection tab.

4. Looking at the protection topology graph, you see an error icon associated with the first volume in the topology, which is the source volume for the SnapMirror relationship.

You also see the horizontal bars in the source volume icon, indicating the warning and error thresholds set for that volume.



- 5. You place your cursor over the error icon to see the pop-up dialog box that displays the threshold settings and see that the volume has exceeded the error threshold, indicating a capacity issue.
- 6. Click the Capacity tab.

Capacity information about volume cluster2 src vol2 displays.

- 7. In the **Capacity** pane, you see that there is an error icon in the bar graph, again indicating that the volume capacity has surpassed the threshold level set for the volume.
- 8. Below the capacity graph, you see that volume autogrow has been disabled and that a volume space guarantee has been set.

You could decide to enable autogrow, but for the purposes of this scenario, you decide to investigate further before making a decision about how to resolve the capacity problem.

- 9. You scroll down to the **Events** list and see that Protection Job Failed, Volume Days Until Full, and Volume Space Full events were generated.
- 10. In the **Events** list, you click the **Volume Space Full** event to get more information, having decided that this event seems most relevant to your capacity issue.

The Event details page displays the Volume Space Full event for the source volume.

- 11. In the **Summary** area, you read the Cause field for the event: The full threshold set at 90% is breached. 45.38 MB (95.54%) of 47.50 MB is used.
- 12. Below the **Summary** area, you see Suggested Corrective Actions.



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.



You can click the name of the aggregate to get further details about the aggregate.

You determine that the aggregate has sufficient space to enable volume autogrow.

b. At the top of the page, look at the icon indicating a critical incident and review the text below the icon.

You determine that "Days to Full: Less than a day | Daily Growth Rate: 5.4%".

14. Go to System Manager or access the ONTAP CLI to enable the volume autogrow option.



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.

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

ManagerDashboards/Overview page to see if there are any problems with your protection relationships and, if they exist, to find solutions.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

In the Dashboards/Overview page, you look at the Unresolved Incidents and Risks area and see a SnapMirror Lag error in the Protection pane under Protection Risks.

Steps

1. In the **Protection** pane on the **Dashboards/Overview** page, locate the SnapMirror relationship lag error and click it.

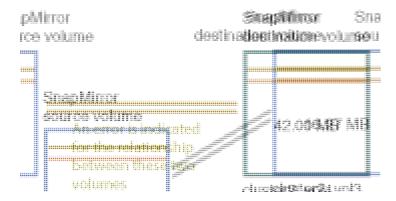
The Event details page for the lag error event is displayed.

- 2. From the **Event** details page you can perform one or more of the following tasks:
 - Review the error message in the Cause field of the Summary area to determine if there is any suggested corrective action.
 - Click the object name, in this case a volume, in the Source field of the Summary area to get details about the volume.
 - Look for notes that might have been added about this event.
 - Add a note to the event.
 - Assign the event to a specific user.
 - Acknowledge or resolve the event.
- 3. In this scenario, you click the object name (in this case, a volume) in the Source field of the **Summary** area to get details about the volume.

The Protection tab of the Health/Volume details page is displayed.

4. In the **Protection** tab, you look at the topology diagram.

You note that the volume with the lag error is the last volume in a three-volume SnapMirror cascade. The volume you selected is outlined in dark gray, and a double orange line from the source volume indicates a SnapMirror relationship error.



5. Click each of the volumes in the SnapMirror cascade.

As you select each volume, the protection information in the Summary, Topology, History, Events, Related Devices, and Related Alerts areas changes to display details relevant to the selected volume.

6. You look at the **Summary** area and position your cursor over the information icon in the **Update Schedule** field for each volume.

In this scenario, you note that the SnapMirror policy is DPDefault, and the SnapMirror schedule updates hourly at five minutes after the hour. You realize that all of the volumes in the relationship are attempting to complete a SnapMirror transfer at the same time.

7. To resolve the lag issue, you modify the schedules for two of the cascaded volumes so that each destination begins a SnapMirror transfer after its source has completed a transfer.

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.

Steps

- 1. Restore data using one of the following tasks:
 - Restore data from the Health/Volume details page.
 - Restore data from the Health/Volumes page.

Restoring data using the Health/Volume details page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Health/Volume details page.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot restore NTFS file streams.

The restore option is not available when:

- The volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered.
- The volume is a FlexGroup volume.
- The volume is configured for SnapMirror Synchronous replication.

Steps

- 1. In the **Protection** tab of the **Health/Volume** details page, right-click in the topology view the name of the volume that you want to restore.
- 2. Select **Restore** from the menu.

Alternatively, select **Restore** from the **Actions** menu to protect the current volume for which you are viewing the details.

The Restore dialog box is displayed.

- 3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
- 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:
 - i. Select the cluster, SVM, and volume to which you want to restore.
 - ii. In the Name table, select a directory name.
 - iii. Click Select Directory.

Click Restore.

The restore process begins.



If a restore operation fails between Cloud Volumes ONTAP HA clusters with an NDMP error, you may need to add an explicit AWS route in the destination cluster so that the destination can communicate with the source system's cluster management LIF. You perform this configuration step using OnCommand Cloud Manager.

Restoring data using the Health/Volumes inventory page

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Health/Volumes inventory page.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot restore NTFS file streams.

The restore option is not available when:

- The volume ID is unknown: for example, when you have a intercluster relationship and the destination cluster has not yet been discovered.
- The volume is a FlexGroup volume.
- The volume is configured for SnapMirror Synchronous replication.

Steps

- 1. In the **Health/Volumes** inventory page, select a volume from which you want to restore data.
- 2. From the toolbar, click **Restore**.

The Restore dialog box is displayed.

- 3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
- 4. Select the items you want to restore.

You can restore the entire volume, or you can specify folders and files you want to restore.

- 5. Select the location to which you want the selected items restored; either **Original Location** or **Alternate Location**.
- 6. Click Restore.

The restore process begins.

Managing health thresholds

You can configure global health threshold values for all the aggregates, volumes, and gtrees to track any health threshold breaches.

What storage capacity health thresholds are

A storage capacity health threshold is the point at which the Unified Manager server generates events to report any capacity problem with storage objects. You can configure alerts to send notification whenever such events occurs.

The storage capacity health thresholds for all aggregates, volumes, and qtrees are set to default values. You can change the settings as required for an object or a group of objects.

Configuring global health threshold settings

You can configure global health threshold conditions for capacity, growth, Snapshot reserve, quotas, and inodes to monitor your aggregate, volume, and qtree size effectively. You can also edit the settings for generating events for exceeding lag thresholds.

About this task

Global health threshold settings apply to all objects with which they are associated, such as aggregates, volumes, and so forth. When thresholds are crossed, an event is generated and, if alerts are configured, an alert notification is sent. Threshold defaults are set to recommended values, but you can modify them to generate events at intervals to meet your specific needs. When thresholds are changed, events are generated or obsoleted in the next monitoring cycle.

Global health threshold settings are accessible from the Configuration/Health Thresholds page. You can also modify threshold settings for individual objects, from the inventory page or the details page for that object.

Choices

Configuring global aggregate health threshold values

You can configure the health threshold settings for capacity, growth, and Snapshot copies for all aggregates to track any threshold breach.

• Configuring global volume health threshold values

You can edit the health threshold settings for capacity, Snapshot copies, qtree quotas, volume growth, overwrite reserve space, and inodes for all volumes to track any threshold breach.

Configuring global qtree health threshold values

You can edit the health threshold settings for capacity for all qtrees to track any threshold breach.

· Editing lag health threshold settings for unmanaged protection relationships

You can increase or decrease the warning or error lag time percentage so that events are generated at intervals that are more appropriate to your needs.

Configuring global aggregate health threshold values

You can configure global health threshold values for all aggregates to track any threshold breach. Appropriate events are generated for threshold breaches and you can take preventive measures based on these events. You can configure the global values based on the best practice settings for thresholds that apply to all monitored aggregates.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you configure the options globally, the default values of the objects are modified. However, if the default values have been changed at the object level, the global values are not modified.

The threshold options have default values for better monitoring, however, you can change the values to suit the requirements of your environment.

When Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.



Health threshold values are not applicable to the root aggregate of the node.

Steps

- 1. In the left navigation pane, click **Configuration > Health Thresholds**.
- 2. In the Configuration/Health Thresholds page, click Aggregates.
- 3. Configure the appropriate threshold values for capacity, growth, and Snapshot copies.
- 4. Click Save.

Configuring global volume health threshold values

You can configure the global health threshold values for all volumes to track any threshold breach. Appropriate events are generated for health threshold breaches, and you can take preventive measures based on these events. You can configure the global values based on the best practice settings for thresholds that apply to all monitored volumes.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Most of the threshold options have default values for better monitoring. However, you can change the values to suit the requirements of your environment.

Note that when Autogrow is enabled on a volume that capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Steps

- 1. In the left navigation pane, click **Configuration > Health Thresholds**.
- 2. In the Configuration/Health Thresholds page, click Volumes.
- 3. Configure the appropriate threshold values for capacity, Snapshot copies, qtree quotas, volume growth, and inodes.
- 4. Click Save.

Configuring global qtree health threshold values

You can configure the global health threshold values for all qtrees to track any threshold breach. Appropriate events are generated for health threshold breaches, and you can take preventive measures based on these events. You can configure the global values based on the best practice settings for thresholds that apply to all monitored gtrees.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The threshold options have default values for better monitoring, however, you can change the values to suit the requirements of your environment.

Events are generated for a qtree only when a Qtree quota or a Default quota has been set on the qtree. Events are not generated if the space defined in a User quota or Group quota has exceeded the threshold.

Steps

- 1. In the left navigation pane, click **Configuration > Health Thresholds**.
- In the Configuration/Health Thresholds page, click Qtrees.
- 3. Configure the appropriate capacity threshold values.
- 4. Click Save.

Editing lag health threshold settings for unmanaged protection relationships

You can edit the global default lag warning and error health threshold settings for unmanaged protection relationships so that events are generated at intervals that are appropriate to your needs.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The lag time must be no more than the defined transfer schedule interval. For example, if the transfer schedule is hourly, then the lag time must not be more than one hour. The lag threshold specifies a percentage that the lag time must not exceed. Using the example of one hour, if the lag threshold is defined as 150%, then you will receive an event when the lag time is more than 1.5 hours.

The settings described in this task are applied globally to all unmanaged protection relationships. The settings cannot be specified and applied exclusively to one unmanaged protection relationship.

Steps

- 1. In the left navigation pane, click **Configuration > Health Thresholds**.
- 2. In the Configuration/Health Thresholds page, click Relationships.
- 3. Increase or decrease the global default warning or error lag time percentage as required.
- 4. Click Save.

Editing individual aggregate health threshold settings

You can edit the health threshold settings for aggregate capacity, growth, and Snapshot copies of one or more aggregates. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

When Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Steps

- 1. In the left navigation pane, click **Health > Aggregates**.
- 2. In the **Health/Aggregates** inventory page, select one or more aggregates and then click **Edit Thresholds**.
- 3. In the **Edit Aggregate Thresholds** dialog box, edit the threshold settings of one of the following: capacity, growth, or Snapshot copies by selecting the appropriate check box and then modifying the settings.
- 4. Click Save.

Editing individual volume health threshold settings

You can edit the health threshold settings for volume capacity, growth, quota, and space

reserve of one or more volumes. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

Note that when Autogrow is enabled on a volume that capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Steps

- 1. In the left navigation pane, click **Health > Volumes**.
- 2. In the Health/Volumes inventory page, select one or more volumes and then click Edit Thresholds.
- In the Edit Volume Thresholds dialog box, edit the threshold settings of one of the following: capacity, Snapshot copies, qtree quota, growth, or inodes by selecting the appropriate check box and then modifying the settings.
- 4. Click Save.

Editing individual qtree health threshold settings

You can edit the health threshold settings for qtree capacity for one or more qtrees. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the **Health/Storage Virtual Machines** inventory page, select the SVM on which the qtree resides.
- 3. In the Health/Storage Virtual Machine details page, click the Qtrees tab.
- 4. Select one or more gtrees and then click Edit Thresholds.
- 5. In the **Edit Qtree Thresholds** dialog box, change the capacity thresholds for the selected qtree or qtrees and click **Save**.

Managing scripts

You can use scripts to automatically modify or update multiple storage objects in Unified

Manager. The script is associated with an alert. When an event triggers an alert, the script is executed. You can upload custom scripts and test their execution when an alert is generated.

How scripts work with alerts

You can associate an alert with your script so that the script is executed when an alert is raised for an event in Unified Manager. You can use the scripts to resolve issues with storage objects or identify which storage objects are generating the events.

When an alert is generated for an event in Unified Manager, an alert email is sent to the specified recipients. If you have associated an alert with a script, the script is executed. You can get the details of the arguments passed to the script from the alert email.

The script uses the following arguments for execution:

- -eventID
- -eventName
- -eventSeverity
- -eventSourceID
- -eventSourceName
- -eventSourceType
- -eventState
- -eventArgs

You can use the arguments in your scripts and gather related event information or modify storage objects.

Example for obtaining arguments from scripts

```
print "$ARGV[0] : $ARGV[1]\n"
print "$ARGV[7] : $ARGV[8]\n"
```

When an alert is generated, this script is executed and the following output is displayed:

```
-eventID : 290
-eventSourceID : 4138
```

Adding scripts

You can add scripts in Unified Manager, and associate the scripts with alerts. These scripts are executed automatically when an alert is generated, and enable you to obtain information about storage objects for which the event is generated.

Before you begin

- · You must have created and saved the scripts that you want to add to the Unified Manager server.
- The supported file formats for scripts are Perl, Shell, PowerShell, and .bat files.
 - For Perl scripts, Perl must be installed on the Unified Manager server. If Perl was installed after Unified Manager, you must restart the Unified Manager server.
 - For PowerShell scripts, the appropriate PowerShell execution policy must be set on the server so that the scripts can be executed.



If your script creates log files to track the alert script progress, you must make sure that the log files are not created anywhere within the Unified Manager installation folder.

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can upload custom scripts and gather event details about the alert.

Steps

- 1. In the toolbar, click , and then click **Scripts** in the left Management menu.
- 2. In the Management/Scripts page, click Add.
- 3. In the Add Script dialog box, click Browse to select your script file.
- 4. Enter a description for the script that you select.
- 5. Click Add.

Deleting scripts

You can delete a script from Unified Manager when the script is no longer required or valid.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The script must not be associated with an alert.

Steps

- 1. In the toolbar, click , and then click **Scripts** in the left Management menu.
- 2. In the Management/Scripts page, select the script that you want to delete, and then click Delete.
- 3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Testing script execution

You can verify that your script is executed correctly when an alert is generated for a storage object.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- You must have uploaded a script in the supported file format to Unified Manager.

Steps

- 1. In the toolbar, click , and then click **Scripts** in the left Management menu.
- 2. In the **Management/Scripts** page page, add your test script.
- 3. In the **Configuration/Alerting** page, perform one of the following actions:

То	Do this
Add an alert	a. In the Configuration/Alerting page, click Add .
	b. In the Actions section, associate the alert with your test script.
Edit an alert	a. In the Configuration/Alerting page, select an alert, and then click Edit .
	b. In the Actions section, associate the alert with your test script.

- 4. Click Save.
- 5. In the Configuration/Alerting page, select the alert that you added or modified, and then click Test.

The script is executed with the "-test" argument, and a notification alert is sent to the email addresses that were specified when the alert was created.

Managing and monitoring groups

You can create groups in Unified Manager to manage storage objects.

Understanding groups

You can create groups in Unified Manager to manage storage objects. Understanding the concepts about groups and how group rules enable you to add storage objects to a group will help you to manage the storage objects in your environment.

What a group is

A group is a dynamic collection of heterogenous storage objects (clusters, SVMs, or volumes). You can create groups in Unified Manager to easily manage a set of storage objects. The members in a group might change, depending on the storage objects that are monitored by Unified Manager at a point in time.

- Each group has a unique name.
- You must configure a minimum of one group rule for each group.

- · You can associate a group with more than one group rule.
- Each group can include multiple types of storage objects such as clusters, SVMs, or volumes.
- Storage objects are dynamically added to a group based on when a group rule is created or when Unified Manager completes a monitoring cycle.
- You can simultaneously apply actions on all the storage objects in a group such as setting thresholds for volumes.

How group rules work for groups

A group rule is a criterion that you define to enable storage objects (volumes, clusters, or SVMs) to be included in a specific group. You can use condition groups or conditions for defining group rule for a group.

- · You must associate a group rule to a group.
- You must associate an object type for a group rule; only one object type is associated for a group rule.
- Storage objects are added or removed from the group after each monitoring cycle or when a rule is created, edited, or deleted.
- A group rule can have one or more condition groups, and each condition group can have one or more conditions.
- Storage objects can belong to multiple groups based on group rules you create.

Conditions

You can create multiple condition groups, and each condition group can have one or more conditions. You can apply all the defined condition groups in a group rule for groups in order to specify which storage objects are included in the group.

Conditions within a condition group are executed using logical AND. All the conditions in a condition group must be met. When you create or modify a group rule, a condition is created that applies, selects, and groups only those storage objects that satisfy all conditions in the condition group. You can use multiple conditions within a condition group when you want to narrow the scope of which storage objects to include in a group.

You can create conditions with storage objects by using the following operands and operator and specifying the required value.

Storage object type	Applicable operands
Volume	Object nameOwning cluster nameOwning SVM nameAnnotations
SVM	Object nameOwning cluster nameAnnotations

Storage object type	Applicable operands
Cluster	Object name
	Annotations

When you select annotation as an operand for any storage object, the "Is" operator is available. For all other operands, you can select either "Is" or "Contains" as operator.

Operand

The list of operands in Unified Manager changes based on the selected object type. The list includes the object name, owning cluster name, owning SVM name, and annotations that you define in Unified Manager.

Operator

The list of operators changes based on the selected operand for a condition. The operators supported in Unified Manager are "Is" and "Contains".

When you select the "Is" operator, the condition is evaluated for exact match of operand value to the value provided for the selected operand.

When you select the "Contains" operator, the condition is evaluated to meet one of the following criteria:

- The operand value is an exact match to the value provided for the selected operand
- The operand value contains the value provided for the selected operand
- Value

The value field changes based on the operand selected.

Example of a group rule with conditions

Consider a condition group for a volume with the following two conditions:

- · Name contains "vol"
- SVM name is "data_svm"

This condition group selects all volumes that include "vol" in their names and that are hosted on SVMs with the name "data_svm".

Condition groups

Condition groups are executed using logical OR, and then applied to storage objects. The storage objects must satisfy one of the condition groups to be included in a group. The storage objects of all the condition groups are combined. You can use condition groups to increase the scope of storage objects to include in a group.

Example of a group rule with condition groups

Consider two condition groups for a volume, with each group containing the following two conditions:

· Condition group 1

- Name contains "vol"
- SVM name is "data_svm"
 Condition group 1 selects all volumes that include "vol" in their names and that are hosted on SVMs with the name "data svm".
- Condition group 2
 - Name contains "vol"
 - The annotation value of data-priority is "critical"
 Condition group 2 selects all volumes that include "vol" in their names and that are annotated with the data-priority annotation value as "critical".

When a group rule containing these two condition groups is applied on storage objects, then the following storage objects are added to a selected group:

- All volumes that include "vol" in their names and that are hosted on the SVM with the name "data svm".
- All volumes that include "vol" in their names and that are annotated with the data-priority annotation value "critical".

How group actions work on storage objects

A group action is an operation that is performed on all the storage objects in a group. For example, you can configure volume threshold group action to simultaneously change the volume threshold values of all volumes in a group.

Groups support unique group action types. You can have a group with only one volume health threshold group action type. However, you can configure a different type of group action, if available, for the same group. The rank of a group action determines the order in which the action is applied to storage objects. The details page of a storage object provides information about which group action is applied on the storage object.

Example of unique group actions

Consider a volume A that belongs to groups G1 and G2, and the following volume health threshold group actions are configured for these groups:

- Change capacity threshold group action with rank 1, for configuring the capacity of the volume
- Change snapshot copies group action with rank 2, for configuring the Snapshot copies of the volume

The <code>Change_capacity_threshold</code> group action always takes priority over the <code>Change_snapshot_copies</code> group action and is applied to volume A. When Unified Manager completes one cycle of monitoring, the health threshold related events of volume A are re-evaluated per the <code>Change_capacity_threshold</code> group action. You cannot configure another volume threshold type of group action for either G1 or G2 group.

Adding groups

You can create groups to combine clusters, volumes, and storage virtual machines (SVMs) for ease of management.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can define group rules to add or remove members from the group and to modify group actions for the group.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- 2. In the **Groups** tab, click **Add**.
- 3. In the Add Group dialog box, enter a name and description for the group.

The group name must be unique.

4. Click Add**.

Editing groups

You can edit the name and description of a group that you created in Unified Manager.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you edit a group to update the name, you must specify a unique name; you cannot use an existing group name.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- 2. In the Groups tab, select the group that you want to edit, and then click Edit.
- 3. In the **Edit Group** dialog box, change the name, description, or both for the group.
- 4. Click Save.

Deleting groups

You can delete a group from Unified Manager when the group is no longer required.

Before you begin

- None of the storage objects (clusters, SVMs, or volumes) must be associated with any group rule that is associated with the group that you want to delete.
- You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click **Management** > **Groups**.
- In the Groups tab, select the group that you want to delete, and then click Delete.
- In the Warning dialog box, confirm the deletion by clicking Yes.

Deleting a group does not delete the group actions that are associated with the group. However, these group actions will be unmapped after the group is deleted.

Adding group rules

You can create group rules for a group to dynamically add storage objects such as volumes, clusters, or storage virtual machines (SVMs) to the group. You must configure at least one condition group with at least one condition to create a group rule.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Storage objects that are currently monitored are added as soon as the group rule is created. New objects are added only after the monitoring cycle is completed.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- 2. In the Group Rules tab, click Add.
- 3. In the Add Group Rule dialog box, specify a name for the group rule.
- 4. In the **Target Object Type** field, select the type of storage object that you want to group.
- 5. In the **Group** field, select the required group for which you want to create group rules.
- 6. In the **Conditions** section, perform the following steps to create a condition, a condition group, or both:

To create	Do this
A condition	a. Select an operand from the list of operands.
	b. Select either Contains or Is as the operator.
	c. Enter a value, or select a value from the available list.
A condition group	a. Click Add Condition Group
	b. Select an operand from the list of operands.
	c. Select either Contains or Is as the operator.
	d. Enter a value, or select a value from the available list.
	e. Click Add condition to create more conditions if required, and repeat steps a through d for each condition.

7. Click Add.

Example for creating a group rule

Perform the following steps in the Add Group Rule dialog box to create a group rule, including configuring a condition and adding a condition group:

- 1. Specify a name for the group rule.
- Select the object type as storage virtual machine (SVM).
- 3. Select a group from the list of groups.
- 4. In the Conditions section, select **Object Name** as the operand.
- 5. Select **Contains** as the operator.
- 6. Enter the value as svm_data.
- 7. Click Add condition group.
- 8. Select **Object Name** as the operand.
- 9. Select **Contains** as the operator.
- 10. Enter the value as vol.
- 11. Click Add condition.
- 12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **critical** as the value in step 10.
- 13. Click **Add** to create the condition for the group rule.

Editing group rules

You can edit group rules to modify the condition groups and the conditions within a condition group to add or remove storage objects to or from a specific group.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- In the toolbar, click , and then click Management > Groups.
- 2. In the Group Rules tab, select the group rule that you want to edit, and then click Edit.
- 3. In the **Edit Group Rule** dialog box, change the group rule name, associated group name, condition groups, and conditions as required.



You cannot change the target object type for a group rule.

Click Save.

Deleting group rules

You can delete a group rule from OnCommand Unified Manager when the group rule is no longer required.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When a group rule is deleted, the associated storage objects will be removed from the group.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- 2. In the **Group Rules** tab, select the group rule that you want to delete, and then click **Delete**.
- 3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Adding group actions

You can configure group actions that you want to apply to storage objects in a group. Configuring actions for a group enables you to save time, because you do not have to add these actions to each object individually.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- 2. In the Group Actions tab, click Add.
- In the Add Group Action dialog box, enter a name and description for the action.
- From the Group menu, select a group for which you want to configure the action.
- 5. From the **Action Type** menu, select an action type.

The dialog box expands, enabling you to configure the selected action type with required parameters.

- Enter appropriate values for the required parameters to configure a group action.
- 7. Click Add.

Editing group actions

You can edit the group action parameters that you configured in Unified Manager, such as the group action name, description, associated group name, and parameters of the action type.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. In the toolbar, click , and then click Management > Groups.

- 2. In the Group Actions tab, select the group action that you want to edit, and then click Edit.
- 3. In the **Edit Group Action** dialog box, change the group action name, description, associated group name, and parameters of the action type, as required.
- 4. Click Save.

Configuring volume health thresholds for groups

You can configure group-level volume health thresholds for capacity, Snapshot copies, qtree quotas, growth, and inodes.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The volume health threshold type of group action is applied only on volumes of a group.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- In the Group Actions tab, click Add.
- 3. Enter a name and description for the group action.
- 4. From the **Group** drop-down box, select a group for which you want to configure group action.
- 5. Select **Action Type** as the volume health threshold.
- 6. Select the category for which you want to set the threshold.
- 7. Enter the required values for the health threshold.
- 8. Click Add.

Deleting group actions

You can delete a group action from Unified Manager when the group action is no longer required.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you delete the group action for the volume health threshold, global thresholds are applied to the storage objects in that group. Any object-level health thresholds that are set on the storage object are not impacted.

Steps

- 1. In the toolbar, click [], and then click Management > Groups.
- 2. In the **Group Actions** tab, select the group action that you want to delete, and then click **Delete**.
- 3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Reordering group actions

You can change the order of the group actions that are to be applied to the storage objects in a group. Group actions are applied to storage objects sequentially based on their rank. The lowest rank is assigned to the group action that you configured last. You can change the rank of the group action depending on your requirements.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can select either a single row or multiple rows, and then perform multiple drag-and-drop operations to change the rank of group actions. However, you must save the changes for the re-prioritization to be reflected in the group actions grid.

Steps

- 1. In the toolbar, click , and then click Management > Groups.
- 2. In the Group Actions tab, click Reorder.
- 3. In the **Reorder Group Actions** dialog box, drag and drop the rows to rearrange the sequence of group actions as required.
- 4. Click Save.

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.

Understanding more about annotations

Understanding the concepts about annotations helps you to manage the events related to the storage objects in your environment.

What annotations are

An annotation is a text string (the name) that is assigned to another text string (the value). Each annotation name-value pair can be dynamically associated with storage objects using annotation rules. When you associate storage objects with predefined annotations, you can filter and view the events that are related to them. You can apply annotations to clusters, volumes, and storage virtual machines (SVMs).

Each annotation name can have multiple values; each name-value pair can be associated with a storage object through rules.

For example, you can create an annotation named "data-center" with the values "Boston" and "Canada". You can then apply the annotation "data-center" with the value "Boston" to volume v1. When an alert is generated for any event on a volume v1 that is annotated with "data-center", the generated email indicates the location of the volume, "Boston", and this enables you to prioritize and resolve the issue.

How annotation rules work in Unified Manager

An annotation rule is a criterion that you define to annotate storage objects (volumes, clusters, or storage virtual machines (SVMs)). You can use either condition groups or conditions for defining annotation rules.

- You must associate an annotation rule to an annotation.
- You must associate an object type for an annotation rule; only one object type can be associated for an annotation rule.
- Unified Manager adds or removes annotations from storage objects after each monitoring cycle or when a rule is created, edited, deleted, or reordered.
- An annotation rule can have one or more condition groups, and each condition group can have one or more conditions.
- Storage objects can have multiple annotations. An annotation rule for a particular annotation can also use different annotations in the rule conditions to add another annotation to already annotated objects.

Conditions

You can create multiple condition groups, and each condition group can have one or more conditions. You can apply all the defined condition groups in an annotation rule of an annotation in order to annotate storage objects.

Conditions within a condition group are executed using logical AND. All the conditions in a condition group must be met. When you create or modify an annotation rule, a condition is created that applies, selects, and annotates only those storage objects that meet all the conditions in the condition group. You can use multiple conditions within a condition group when you want to narrow the scope of which storage objects to annotate.

You can create conditions with storage objects by using the following operands and operator and specifying the required value.

Storage object type	Applicable operands
Volume	Object nameOwning cluster nameOwning SVM name
SVM	AnnotationsObject nameOwning cluster nameAnnotations
Cluster	Object name Annotations

When you select annotation as an operand for any storage object, the "Is" operator is available. For all other operands, you can select either "Is" or "Contains" as operator. When you select the "Is" operator, the condition is evaluated for an exact match of the operand value with the value provided for the selected operand. When you select the "Contains" operator, the condition is evaluated to meet one of the following criteria:

- The operand value is an exact match to the value of the selected operand.
- The operand value contains the value provided for the selected operand.

Example of an annotation rule with conditions

Consider an annotation rule with one condition group for a volume with the following two conditions:

- · Name contains "vol"
- SVM name is "data_svm"

This annotation rule annotates all volumes that include "vol" in their names and that are hosted on SVMs with the name "data_svm" with the selected annotation and the annotation type.

Condition groups

Condition groups are executed using logical OR, and then applied to storage objects. The storage objects must meet the requirements of one of the condition groups to be annotated. The storage objects that meet the conditions of all the condition groups are annotated. You can use condition groups to increase the scope of storage objects to be annotated.

Example of an annotation rule with condition groups

Consider an annotation rule with two condition groups for a volume; each group contains the following two conditions:

- · Condition group 1
 - Name contains "vol"
 - SVM name is "data_svm"
 This condition group annotates all volumes that include "vol" in their names and that are hosted on SVMs with the name "data svm".
- Condition group 2
 - Name contains "vol"
 - The annotation value of data-priority is "critical"
 This condition group annotates all volumes that include "vol" in their names and that are annotated with the data-priority annotation value as "critical".

When an annotation rule containing these two condition groups is applied on storage objects, then the following storage objects are annotated:

- · All volumes that include "vol" in their names and that are hosted on SVM with the name "data svm".
- All volumes that include "vol" in their names and that are annotated with the data-priority annotation value as "critical".

Description of predefined annotation values

Data-priority is a predefined annotation that has the values Mission critical, High, and Low. These values enable you to annotate storage objects based on the priority of data that they contain. You cannot edit or delete the predefined annotation values.

· Data-priority: Mission critical

This annotation is applied to storage objects that contain mission-critical data. For example, objects that contain production applications can be considered as mission critical.

· Data-priority:High

This annotation is applied to storage objects that contain high-priority data. For example, objects that are hosting business applications can be considered high priority.

Data-priority:Low

This annotation is applied to storage objects that contain low-priority data. For example, objects that are on secondary storage, such as backup and mirror destinations, might be of low priority.

Adding annotations dynamically

When you create custom annotations, Unified Manager dynamically associates clusters, storage virtual machines (SVMs), and volumes with the annotations by using rules. These rules automatically assign the annotations to storage objects.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the Annotations page, click Add Annotation.
- 3. In the **Add Annotation** dialog box, type a name and description for the annotation.

You can also add values to annotations while creating annotations.

- 4. Optional: In the Annotation Values section, click Add to add values to the annotation.
- Click Save and Close.

Adding values to annotations

You can add values to annotations, and then associate storage objects with a particular annotation name-value pair. Adding values to annotations helps you to manage storage objects more effectively.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You cannot add values to predefined annotations.

Steps

1. In the toolbar, click , and then click **Annotations** in the left Management menu.

- 2. In the **Annotations** page, select the annotation to which you want to add a value and then click **Add** in the **Values** section.
- 3. In the Add Annotation Value dialog box, specify a value for the annotation.

The value that you specify must be unique for the selected annotation.

Click Add.

Deleting annotations

You can delete custom annotations and their values when they are no longer required.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The annotation values must not be used in other annotations or group rules.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the **Annotations** tab, select the annotation that you want to delete.

The details of the selected annotation are displayed.

- Click Actions > Delete to delete the selected annotation and its value.
- 4. In the warning dialog box, click **Yes** to confirm the deletion.

Results

The selected annotation and its value is deleted.

Viewing the annotation list and details

You can view the list of annotations that are dynamically associated with clusters, volumes, and storage virtual machines (SVMs). You can also view details such as the description, created by, created date, values, rules, and the objects associated with the annotation.

Steps

- 1. In the toolbar, click (), and then click **Annotations** in the left Management menu.
- 2. In the Annotations tab, click the annotation name to view the associated details.

Deleting values from annotations

You can delete values associated with custom annotations when that value no longer applies to the annotation.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

• The annotation value must not be associated with any annotation rules or group rules.

About this task

You cannot delete values from predefined annotations.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the annotations list in the **Annotations** tab, select the annotation from which you want to delete a value.
- 3. In the Values area of the Annotations tab, select the value you want to delete, and then click Delete.
- In the Warning dialog box, click Yes.

The value is deleted and no longer displayed in the list of values for the selected annotation.

Creating annotation rules

You can create annotation rules that Unified Manager uses to dynamically annotate storage objects such as volumes, clusters, or storage virtual machines (SVMs).

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Storage objects that are currently monitored are annotated as soon as the annotation rule is created. New objects are annotated only after the monitoring cycle is completed.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the Annotation Rules tab, click Add.
- 3. In the Add Annotation Rule dialog box, specify a name for the annotation rule.
- 4. In the **Target Object Type** field, select the type of storage object that you want to annotate.
- 5. In the Apply Annotation fields, select the annotation and annotation value that you want to use.
- 6. In the **Conditions** section, perform the appropriate action to create a condition, a condition group, or both:

To create	Do this
A condition	a. Select an operand from the list of operands.b. Select either Contains or Is as the operator.
	c. Enter a value, or select a value from the available list.

To create	Do this
A condition group	a. Click Add Condition Group.
	b. Select an operand from the list of operands.
	c. Select either Contains or Is as the operator.
	d. Enter a value, or select a value from the available list.
	e. Click Add condition to create more conditions if required, and repeat steps a through d for each condition.

7. Click Add.

Example of creating an annotation rule

Perform the following steps in the Add Annotation Rule dialog box to create an annotation rule, including configuring a condition and adding a condition group:

- 1. Specify a name for the annotation rule.
- 2. Select the target object type as storage virtual machine (SVM).
- 3. Select an annotation from the list of annotations, and specify a value.
- 4. In the Conditions section, select **Object Name** as the operand.
- 5. Select **Contains** as the operator.
- 6. Enter the value as svm data.
- 7. Click Add condition group.
- 8. Select **Object Name** as the operand.
- 9. Select **Contains** as the operator.
- 10. Enter the value as vol.
- 11. Click Add condition.
- 12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **mission-critical** as the value in step 10.
- 13. Click Add.

Adding annotations manually to individual storage objects

You can manually annotate selected volumes, clusters, and SVMs without using annotation rules. You can annotate a single storage object or multiple storage objects, and specify the required name-value pair combination for the annotation.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

1. Navigate to the storage objects you want to annotate:

To add annotation to	Do this
Clusters	a. Click Health > Clusters.b. Select one or more clusters.
Volumes	a. Click Health > Volumes.b. Select one or more volumes.
SVMs	a. Click Health > SVMs.b. Select one or more SVMs.

- 2. Click Annotate and select a name-value pair.
- 3. Click Apply.

Editing annotation rules

You can edit annotation rules to modify the condition groups and conditions within the condition group to add annotations to or remove annotations from storage objects.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Annotations are dissociated from storage objects when you edit the associated annotation rules.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the Annotation Rules tab, select the annotation rule you want to edit, and then click Actions > Edit.
- 3. In the **Edit Annotation Rule** dialog box, change the rule name, annotation name and value, condition groups, and conditions as required.

You cannot change the target object type for an annotation rule.

4. Click Save.

Configuring conditions for annotation rules

You can configure one or more conditions to create annotation rules that Unified Manager applies on the storage objects. The storage objects that satisfy the annotation rule are annotated with the value specified in the rule.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the Annotation Rules tab, click Add.
- 3. In the Add Annotation Rule dialog box, enter a name for the rule.
- 4. Select one object type from the Target Object Type list, and then select an annotation name and value from the list.
- 5. In the **Conditions** section of the dialog box, select an operand and an operator from the list and enter a condition value, or click **Add Condition** to create a new condition.
- 6. Click Save and Add.

Example of configuring a condition for an annotation rule

Consider a condition for the object type SVM, where the object name contains "svm_data".

Perform the following steps in the Add Annotation Rule dialog box to configure the condition:

- 1. Enter a name for the annotation rule.
- 2. Select the target object type as SVM.
- 3. Select an annotation from the list of annotations and a value.
- 4. In the Conditions field, select Object Name as the operand.
- 5. Select **Contains** as the operator.
- 6. Enter the value as svm_data.
- 7. Click Add.

Deleting annotation rules

You can delete annotation rules from OnCommand Unified Manager when the rules are no longer required.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

When you delete an annotation rule, the annotation is disassociated and removed from the storage objects.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the **Annotation Rules** tab, select the annotation rule that you want to delete, and then click **Delete**.
- 3. In the **Warning** dialog box, click **Yes** to confirm the deletion.

Reordering annotation rules

You can change the order in which Unified Manager applies annotation rules to storage objects. Annotation rules are applied to storage objects sequentially based on their rank. When you configure an annotation rule, the rank is least. But you can change the rank of the annotation rule depending on your requirements.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can select either a single row or multiple rows and perform many drag-and-drop operations to change the rank of annotation rules. However, you must save the changes for the reprioritization to be displayed in the Annotation Rules tab.

Steps

- 1. In the toolbar, click , and then click **Annotations** in the left Management menu.
- 2. In the Annotation Rules tab, click Reorder.
- 3. In the **Reorder Annotation Rule** dialog box, drag and drop single or multiple rows to rearrange the sequence of the annotation rules.
- 4. Click Save.

You must save the changes for the reorder to be displayed.

Configuring backup and restore operations

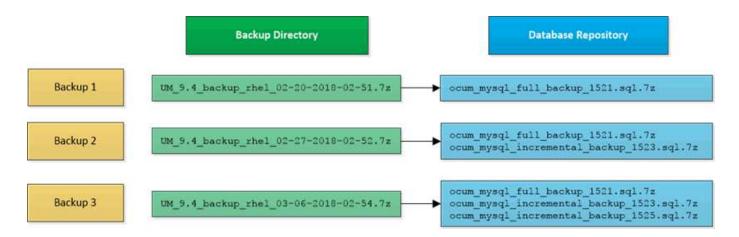
You can create backups of Unified Manager and use the restore feature to restore the backup to the same (local) system or a new (remote) system in case of a system failure or data loss.

What a database backup is

A backup is a copy of the Unified Manager database and configuration files that you can use in case of a system failure or data loss. You can schedule a backup to be written to a local destination or to a remote destination. It is highly recommended that you define a remote location that is external to the Unified Manager host system.

A backup consists of a single file in the backup directory and one or more files in the database repository directory. The file in the backup directory is very small because it contains only a pointer to the files located in the database repository directory that are required to recreate the backup.

The first time you generate a backup a single file is created in the backup directory and a full backup file is created in the database repository directory. The next time you generate a backup a single file is created in the backup directory and an incremental backup file is created in the database repository directory that contains the differences from the full backup file. This process continues as you create additional backups, up to the maximum retention setting, as shown in the following figure.





Do not rename or remove any of the backup files in these two directories or any subsequent restore operation will fail.

If you write your backup files to the local system, you should initiate a process to copy the backup files to a remote location so they will be available in case you have a system issue that requires a complete restore.

Before beginning a backup operation, Unified Manager performs an integrity check to verify that all the required backup files and backup directories exist and are writable. It also checks that there is enough space on the system to create the backup file.

Note that you can restore a backup only on the same version of Unified Manager. For example, if you created a backup on Unified Manager 9.4, the backup can be restored only on Unified Manager 9.4 systems.

Configuring database backup settings

You can configure the Unified Manager database backup settings to set the database backup path, retention count, and backup schedules. You can enable daily or weekly scheduled backups. By default, scheduled backups are disabled.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You must have a minimum of 150 GB of space available in the location you define as the backup path.

It is recommended that you use a remote location that is external to the Unified Manager host system.

- When Unified Manager is installed on a Linux system, verify that the "jboss" user has write permissions to the backup directory.
- You should not schedule backup operations to occur immediately after a new cluster has been added while Unified Manager is collecting 15 days of historical performance data.

About this task

More time is required the first time a backup is performed than for subsequent backups because the first backup is a full backup. A full backup can be over 1 GB and can take three to four hours. Subsequent backups are incremental and require less time.

Steps

- 1. In the toolbar, click , and then click Management > Database Backup.
- 2. In the Management/Database Backup page, click Actions > Database Backup Settings.
- 3. Configure the appropriate values for a backup path and retention count.

The default value for retention count is 10; you can use 0 for creating unlimited backups.

- 4. In the **Schedule Frequency** section, select the **Enable** checkbox, and then specify a daily or weekly schedule.
 - Daily

If you select this option, you must enter a time in 24-hour format for creating the backup. For example, if you specify 18:30, then a backup is created daily at 6:30 PM.

Weekly

If you select this option, you must specify the time and day for creating the backup. For example, if you specify the day as Monday and time as 16:30, then a weekly backup is created every Monday at 4:30 PM.

5. Click Save and Close.

What a database restore is

Database restore is the process of restoring an existing Unified Manager backup file to the same or a different Unified Manager server. You perform the restore operation from the Unified Manager console.

If you are performing a restore operation on the same (local) system, and the backup files are all stored locally, you can run the restore command using the default location. If you are performing a restore operation on a different Unified Manager system (a remote system), you must copy the backup file, or files, from secondary storage to the local disk before running the restore command.

During the restore process, you are logged out of Unified Manager. You can log in to the system after the restore process is complete.

The restore feature is version-specific and platform-specific. You can restore a Unified Manager backup only on the same version of Unified Manager. Unified Manager supports backup and restore in the following platform scenarios:

- · Virtual appliance to virtual appliance
- Virtual appliance to Red Hat Enterprise Linux or CentOS
- Red Hat Enterprise Linux to Red Hat Enterprise Linux or CentOS
- · Windows to Windows

If you are restoring the backup image to a new server, after the restore operation completes you need to generate a new HTTPS security certificate and restart the Unified Manager server. You will also need to reconfigure SAML authentication settings, if they are required, when restoring the backup image to a new server.



Old backup files cannot be used to restore an image after Unified Manager has been upgraded to a newer version of software. To save space, all old backup files, except the newest file, are removed automatically when you upgrade Unified Manager.

Virtual appliance backup and restore process overview

The backup and restore model for Unified Manager when installed on a virtual appliance is to capture and restore an image of the full virtual application.

Because the Unified Manager backup operation on the virtual appliance does not provide a way to move the backup file off of the vApp, the following tasks enable you to complete a backup of the virtual appliance:

- 1. Power off the VM and take a VMware snapshot of the Unified Manager virtual appliance.
- 2. Make a NetApp Snapshot copy on the datastore to capture the VMware snapshot.

If the datastore is not hosted on a system running ONTAP software, follow the storage vendor guidelines to create a backup of the VMware snapshot.

- 3. Replicate the NetApp Snapshot copy, or snapshot equivalent, to alternate storage.
- 4. Delete the VMware snapshot.

You should implement a backup schedule using these tasks to ensure that the Unified Manager virtual appliance is protected if issues arise.

To restore the VM, you can use the VMware snapshot you created to restore the VM to the backup point-in-time state.

Restoring a database backup on a virtual machine

In case of data loss or data corruption, you can use the restore feature to restore Unified Manager to the previous stable state with minimal loss. You can restore the Unified Manager database on a virtual machine by using the Unified Manager maintenance console.

Before you begin

- · You must have the maintenance user credentials.
- The Unified Manager backup files must be on the local system.
- The backup files must be of .7z type.

About this task

Backup compatibility is platform and version dependent. You can restore a backup from a virtual appliance to another virtual appliance, or from a virtual appliance to a Red Hat Enterprise Linux or CentOS system.



When performing a restore operation on a different virtual appliance than the system from which the original backup file was created, the maintenance user name and password on the new vApp must be the same as the credentials from the original vApp.

Steps

- 1. In the vSphere client, locate the Unified Manager virtual machine, and then select the **Console** tab.
- Click in the console window, and then log in to the maintenance console using your user name and password.
- 3. In the **Main Menu**, enter the number for the **System Configuration** option.
- 4. In the System Configuration Menu, enter the number for the Restore from an OCUM Backup option.
- 5. When prompted, enter the absolute path of the backup file.

```
Bundle to restore from: opt/netapp/data/ocum-backup/UM_9.4.N151112.0947_backup_unix_02-25-2018-11-41.7z
```

After the restore operation is complete, you can log in to Unified Manager.

After you finish

After you restore the backup, if the OnCommand Workflow Automation server does not work, perform the following steps:

- 1. On the Workflow Automation server, change the IP address of the Unified Manager server to point to the latest machine.
- 2. On the Unified Manager server, reset the database password if the acquisition fails in step 1.

Restoring a database backup on a Linux system

If data loss or data corruption occurs, you can restore Unified Manager to the previous stable state with minimum loss of data. You can restore the Unified Manager database to a local or remote Red Hat Enterprise Linux or CentOS system.

Before you begin

- You must have Unified Manager installed on a server.
- You must have the root user credentials for the Linux host on which Unified Manager is installed.
- You must have copied the Unified Manager backup file and the contents of the database repository directory to the system on which you will perform the restore operation.

It is recommended that you copy the backup file to the default directory /data/ocum-backup. The database repository files must be copied to the /database-dumps-repo subdirectory under the /ocum-backup directory.

The backup files must be of .7z type.

About this task

The restore feature is platform-specific and version-specific. You can restore a Unified Manager backup only on the same version of Unified Manager. You can restore a Linux backup file or a virtual appliance backup file to a Red Hat Enterprise Linux or CentOS system.



If the backup folder name contains a space, you must include the absolute path or relative path in double quotation marks.

Steps

- 1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
- 2. Log in as the root user to the host on which Unified Manager is installed.
- 3. If Unified Manager is installed in VCS setup, then stop the Unified Manager ocie and ocieau services using Veritas Operations Manager.

After you finish

After the restore operation is complete, you can log in to Unified Manager.

Restoring a database backup on Windows

In case of data loss or data corruption, you can use the restore feature to restore Unified Manager to the previous stable state with minimal loss. You can restore the Unified Manager database to a local Windows system or a remote Windows system by using the restore command.

Before you begin

- · You must have Unified Manager installed on a server.
- · You must have Windows administrator privileges.
- You must have copied the Unified Manager backup file and the contents of the database repository directory to the system on which you will perform the restore operation.

It is recommended that you copy the backup file to the default directory \ProgramData\NetApp\OnCommandAppData\ocum\backup. The database repository files must be copied to the \database dumps repo subdirectory under the \backup directory.

The backup files must be of .7z type.

About this task

The restore feature is platform-specific and version-specific. You can restore a Unified Manager backup only on the same version of Unified Manager, and a Windows backup can be restored only on a Windows platform.



If the folder names contain a space, you must include the absolute path or relative path of the backup file in double quotation marks.

Steps

- 1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
- 2. Log in to the Unified Manager console as an administrator: um cli login -u maint_username

After you finish

After the restore operation is complete, you can log in to Unified Manager.

Migrating a Unified Manager virtual appliance to a Linux system

You can restore a Unified Manager database backup from a virtual appliance to a Red Hat Enterprise Linux or CentOS Linux system if you want to change the host operating system on which Unified Manager is running.

Before you begin

- · On the virtual appliance:
 - You must have the Operator, OnCommand Administrator, or Storage Administrator role to create the backup.
 - You must know the name of the Unified Manager maintenance user for the restore operation.
- On the Linux system:
 - You must have installed Unified Manager on a RHEL or CentOS server following the instructions in the Installation Guide.
 - The version of Unified Manager on this server must be the same as the version on the virtual appliance from which you are using the backup file.
 - Do not launch the UI or configure any clusters, users, or authentication settings on the Linux system after installation. The backup file populates this information during the restore process.
 - You must have the root user credentials for the Linux host.

About this task

These steps describe how to create a backup file on the virtual appliance, copy the backup files to the Red Hat Enterprise Linux or CentOS system, and then restore the database backup to the new system.

Steps

- 1. On the virtual appliance, in the toolbar click ; and then click **Management** > **Database Backup**.
- 2. In the Management/Database Backup page, click Actions > Database Backup Settings.
- 3. Change the backup path to /jail/support.

- 4. In the **Schedule Frequency** section, select the **Enable** checkbox, select **Daily**, and enter a time a few minutes past the current time so that the backup is created shortly.
- 5. Click Save and Close.
- 6. Wait a few hours for the backup to be generated.

A full backup can be over 1 GB and can take three to four hours to complete.

7. Log in as the root user to the Linux host on which Unified Manager is installed and copy the backup files from /support on the virtual appliance using SCP.root@<rhel_server>:/# scp -r admin@<vapp server ip address>:/support/* .

```
root@ocum rhel-21:/# scp -r admin@10.10.10.10:/support/* .
```

Make sure you have copied the .7z backup file and all the .7z repository files in the <code>/database-dumps-repo</code> subdirectory.

8. At the command prompt, restore the backup: um backup restore -f
 /<backup_file_path>/<backup_file_name>
 um backup restore -f /UM 9.4.N151113.1348 backup unix 02-12-2018-04-16.7z

9. After the restore operation completes, log in to the Unified Manager web UI.

After you finish

You should perform the following tasks:

- Generate a new HTTPS security certificate and restart the Unified Manager server.
- Change the backup path to the default setting for your Linux system (/data/ocum-backup), or to a new
 path of your choice, because there is no /jail/support path on the Linux system.
- · Reconfigure both sides of your Workflow Automation connection, if WFA is being used.
- · Reconfigure SAML authentication settings, if you are using SAML.

After you have verified that everything is running as expected on your Linux system, you can shut down and remove the Unified Manager virtual appliance.

What a Unified Manager maintenance window is

You define a Unified Manager maintenance window to suppress events and alerts for a specific timeframe when you have scheduled cluster maintenance and you do not want to receive a flood of unwanted notifications.

When the maintenance window starts, an "Object Maintenance Window Started" event is posted to the Events inventory page. This event is obsoleted automatically when the maintenance window ends.

During a maintenance window the events related to all objects on that cluster are still generated, but they do not appear in any of the UI pages, and no alerts or other types of notification are sent for these events. You can, however, view the events that were generated for all storage objects during a maintenance window by selecting one of the View options on the Events inventory page.

You can schedule a maintenance window to be initiated in the future, you can change the start and end times

for a scheduled maintenance window, and you can cancel a scheduled maintenance window.

Scheduling a maintenance window to disable cluster event notifications

If you have a planned downtime for a cluster, for example, to upgrade the cluster or to move one of the nodes, you can suppress the events and alerts that would normally be generated during that timeframe by scheduling a Unified Manager maintenance window.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

During a maintenance window, the events related to all objects on that cluster are still generated, but they do not appear in the event page, and no alerts or other types of notification are sent for these events.

The time you enter for the maintenance window is based on the time at the Unified Manager server.

Steps

- 1. In the left navigation pane, click Configuration > Cluster Data Sources.
- 2. In the Maintenance Mode column for the cluster, select the slider button and move it to the right.

The calendar window is displayed.

3. Select the start and end date and time for the maintenance window and click Apply.

The message "Scheduled" appears next to the slider button.

Results

When the start time is reached the cluster goes into maintenance mode and an "Object Maintenance Window Started" event is generated.

Changing or canceling a scheduled maintenance window

If you have configured a Unified Manager maintenance window to occur in the future, you can change the start and end times or cancel the maintenance window from occurring.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

Canceling a currently running maintenance window is useful if you have completed cluster maintenance before the scheduled maintenance window end time and you want to start receiving events and alerts from the cluster again.

Steps

1. In the left navigation pane, click **Configuration > Cluster Data Sources**.

2. In the **Maintenance Mode** column for the cluster:

If you want to	Perform this step
Change the timeframe for a scheduled maintenance window	a. Click the text "Scheduled" next to the slider button.
	b. Change the start and/or end date and time and click Apply .
Extend the length of an active maintenance window	a. Click the text "Active" next to the slider button.
	b. Change the end date and time and click Apply .
Cancel a scheduled maintenance window	Select the slider button and move it to the left.
Cancel an active maintenance window	Select the slider button and move it to the left.

Viewing events that occurred during a maintenance window

If necessary, you can view the events that were generated for all storage objects during a Unified Manager maintenance window. Most events will appear in the Obsolete state once the maintenance window has completed and all system resources are back up and running.

Before you begin

At least one maintenance window must have completed before any events are available.

About this task

Events that occurred during a maintenance window do not appear on the Events inventory page by default.

Steps

- 1. In the left navigation pane, click **Events**.
 - By default, all active (New and Acknowledged) events are displayed on the Events inventory page.
- 2. From the View pane, select the option All events generated during maintenance.
 - The list of events trigged during the last 7 days from all maintenance window sessions and from all clusters are displayed.
- 3. If there have been multiple maintenance windows for a single cluster, you can click the **Triggered Time** calendar icon and select the period of time for the maintenance window events that you are interested in viewing.

Managing SAML authentication settings

After you have configured remote authentication settings, you can enable Security Assertion Markup Language (SAML) authentication so that remote users are

authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Note that only remote users will have access to the Unified Manager graphical user interface after SAML authentication has been enabled. Local users and Maintenance users will not be able to access the UI. This configuration does not impact users who access the maintenance console.

Identity provider requirements

When configuring Unified Manager to use an identity provider (IdP) to perform SAML authentication for all remote users, you need to be aware of some required configuration settings so that the connection to Unified Manager is successful.

You must enter the Unified Manager URI and metadata into the IdP server. You can copy this information from the Unified ManagerSAML Authentication page. Unified Manager is considered the service provider (SP) in the Security Assertion Markup Language (SAML) standard.

Supported encryption standards

- Advanced Encryption Standard (AES): AES-128 and AES-256
- Secure Hash Algorithm (SHA): SHA-1 and SHA-256

Validated identity providers

- Shibboleth
- Active Directory Federation Services (ADFS)

ADFS configuration requirements

• You must define three claim rules in the following order that are required for Unified Manager to parse ADFS SAML responses for this relying party trust entry.

Claim rule	Value
SAM-account-name	Name ID
SAM-account-name	urn:oid:0.9.2342.19200300.100.1.1
Token groups — Unqualified Name	urn:oid:1.3.6.1.4.1.5923.1.5.1.1

- You must set the authentication method to "Forms Authentication" or users may receive an error when logging out of Unified Manager when using Internet Explorer. Follow these steps:
 - a. Open the ADFS Management Console.
 - b. Click on the Authentication Policies folder on the left tree view.
 - c. Under Actions on the right, click Edit Global Primary Authentication Policy.
 - d. Set the Intranet Authentication Method to "Forms Authentication" instead of the default "Windows Authentication".
- In some cases login through the IdP is rejected when the Unified Manager security certificate is CA-signed. There are two workarounds to resolve this issue:

 Follow the instructions identified in the link to disable the revocation check on the ADFS server for chained CA cert associated relying party:

http://www.torivar.com/2016/03/22/adfs-3-0-disable-revocation-check-windows-2012-r2/

Have the CA server reside within the ADFS server to sign the Unified Manager server cert request.

Other configuration requirements

- The Unified Manager clock skew is set to 5 minutes, so the time difference between the IdP server and the Unified Manager server cannot be more than 5 minutes or authentication will fail.
- When users attempt to access Unified Manager using Internet Explorer they might see the message The
 website cannot display the page. If this occurs, make sure these users uncheck the option for "Show
 friendly HTTP error messages" in Tools > Internet Options > Advanced.

Enabling SAML authentication

You can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Before you begin

- · You must have configured remote authentication and verified that it is successful.
- You must have created at least one Remote User, or a Remote Group, with the OnCommand Administrator role.
- The Identity provider (IdP) must be supported by Unified Manager and it must be configured.
- You must have the IdP URL and metadata.
- · You must have access to the IdP server.

About this task

After you have enabled SAML authentication from Unified Manager, users cannot access the graphical user interface until the IdP has been configured with the Unified Manager server host information. So you must be prepared to complete both parts of the connection before starting the configuration process. The IdP can be configured before or after configuring Unified Manager.

Only remote users will have access to the Unified Manager graphical user interface after SAML authentication has been enabled. Local users and Maintenance users will not be able to access the UI. This configuration does not impact users who access the maintenance console, the Unified Manager commands, or ZAPIs.



Unified Manager is restarted automatically after you complete the SAML configuration on this page.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, select the SAML Authentication tab.
- Select the Enable SAML authentication checkbox.

The fields required to configure the IdP connection are displayed.

4. Enter the IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP server.

If the IdP server is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URI to populate the IdP Metadata field automatically.

5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.

You can configure the IdP server with this information at this time.

6. Click Save.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

7. Click **Confirm and Logout** and Unified Manager is restarted.

Results

The next time authorized remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the IdP login page instead of the Unified Manager login page.

After you finish

If not already completed, access your IdP and enter the Unified Manager server URI and metadata to complete the configuration.



When using ADFS as your identity provider, the Unified Manager GUI does not honor the ADFS timeout and will continue to work until the Unified Manager session timeout is reached. When Unified Manager is deployed on Windows, Red Hat, or CentOS, you can change the GUI session timeout using the following Unified Manager CLI command: um option set absolute.session.timeout=00:15:00This command sets the Unified Manager GUI session timeout to 15 minutes.

Changing the identity provider used for SAML authentication

You can change the identity provider (IdP) that Unified Manager uses to authenticate remote users.

Before you begin

- You must have the IdP URL and metadata.
- · You must have access to the IdP.

About this task

The new IdP can be configured before or after configuring Unified Manager.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the Setup/Authentication page, select the SAML Authentication tab.
- 3. Enter the new IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP.

If the IdP is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URL to populate the IdP Metadata field automatically.

- 4. Copy the Unified Manager metadata URI, or save the metadata to an XML text file.
- 5. Click Save Configuration.

A message box displays to confirm that you want to change the configuration.

6. Click OK.

After you finish

Access the new IdP and enter the Unified Manager server URI and metadata to complete the configuration.

The next time authorized remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the new IdP login page instead of the old IdP login page.

Updating SAML authentication settings after Unified Manager security certificate change

Any change to the HTTPS security certificate installed on the Unified Manager server requires that you update the SAML authentication configuration settings. The certificate is updated if you rename the host system, assign a new IP address for the host system, or manually change the security certificate for the system.

About this task

After the security certificate is changed and the Unified Manager server is restarted, SAML authentication will not function and users will not be able to access the Unified Manager graphical interface. You must update the SAML authentication settings on both the IdP server and on the Unified Manager server to re-enable access to the user interface.

Steps

- 1. Log into the maintenance console.
- 2. In the Main Menu, enter the number for the Disable SAML authentication option.

A message displays to confirm that you want to disable SAML authentication and restart Unified Manager.

- Launch the Unified Manager user interface using the updated FQDN or IP address, accept the updated server certificate into your browser, and log in using the maintenance user credentials.
- 4. In the Setup/Authentication page, select the SAML Authentication tab and configure the IdP connection.
- 5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.
- 6. Click Save.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

- 7. Click **Confirm and Logout** and Unified Manager is restarted.
- 8. Access your IdP server and enter the Unified Manager server URI and metadata to complete the configuration.

Identity provider	Configuration steps
ADFS	a. Delete the existing relying party trust entry in the ADFS management GUI.
	b. Add a new relying party trust entry using the saml_sp_metadata.xml from the updated Unified Manager server.
	 c. Define the three claim rules that are required for Unified Manager to parse ADFS SAML responses for this relying party trust entry.
	d. Restart the ADFS Windows service.
Shibboleth	a. Update the new FQDN of Unified Manager server into the attribute-filter.xml and relying-party.xml files.
	b. Restart the Apache Tomcat web server and wait for port 8005 to come online.

9. Log in to Unified Manager and verify that SAML authentication works as expected through your IdP.

Disabling SAML authentication

You can disable SAML authentication when you want to stop authenticating remote users through a secure identity provider (IdP) before they can log into the Unified Manager web UI. When SAML authentication is disabled, the configured directory service providers, such as Active Directory or LDAP, perform sign-on authentication.

About this task

After you disable SAML authentication, Local users and Maintenance users will be able to access the graphical user interface in addition to configured Remote users.

You can also disable SAML authentication using the Unified Manager maintenance console if you do not have access to the graphical user interface.



Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

- 1. In the toolbar, click , and then click **Authentication** in the left Setup menu.
- 2. In the **Setup/Authentication** page, select the **SAML Authentication** tab.
- 3. Uncheck the Enable SAML authentication checkbox.
- 4. Click Save.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

5. Click **Confirm and Logout** and Unified Manager is restarted.

Results

The next time remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the Unified Manager login page instead of the IdP login page.

After you finish

Access your IdP and delete the Unified Manager server URI and metadata.

Disabling SAML authentication from the maintenance console

You may need to disable SAML authentication from the maintenance console when there is no access to the Unified Manager GUI. This could happen in cases of mis-configuration or if the IdP is not accessible.

Before you begin

You must have access to the maintenance console as the maintenance user.

About this task

When SAML authentication is disabled, the configured directory service providers, such as Active Directory or LDAP, perform sign-on authentication. Local users and Maintenance users will be able to access the graphical user interface in addition to configured Remote users.

You can also disable SAML authentication from the Setup/Authentication page in the UI.



Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

- 1. Log into the maintenance console.
- 2. In the Main Menu, enter the number for the Disable SAML authentication option.

A message displays to confirm that you want to disable SAML authentication and restart Unified Manager.

3. Type y, and then press Enter and Unified Manager is restarted.

Results

The next time remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the Unified Manager login page instead of the IdP login page.

After you finish

If required, access your IdP and delete the Unified Manager server URL and metadata.

Managing storage objects using the Favorites option

The Favorites option enables you to view and manage selected storage objects in Unified Manager by marking them as favorites. You can quickly view the status of your favorite storage objects and fix issues before they become critical.

Tasks you can perform from the Favorites dashboard

- · View the list of storage objects marked as favorite.
- · Add storage objects to the Favorites list.
- Remove storage objects from the Favorites list.

Viewing the Favorites list

You can view the capacity, performance, and protection details of selected storage objects from the Favorites list. The details of a maximum of 20 storage objects are displayed in the Favorites list.

Adding storage objects to the Favorites list

You can add storage objects to the Favorites list, and then monitor these objects for health, capacity, and performance. You can only mark clusters, volumes, and aggregates as favorite.

Removing storage objects from the Favorites list

You can remove storage objects from the Favorites list when you no longer require them to be marked as favorite.

Adding to, and removing storage objects from, the Favorites list

You can add storage objects to a Favorites list so you can monitor the objects for health, capacity, and performance. You can use object status in the Favorites list to determine issues and fix them before they become critical. The Favorites list also provides the most recent monitoring status of a storage object. You can remove storage objects from the Favorites list when you no longer require them to be marked as favorite.

About this task

You can add up to 20 clusters, nodes, aggregates, or volumes to the Favorites list. When you add a node to the Favorites list, it is displayed as a cluster.

Steps

- 1. Go to the **Details** page of the storage object that you want to mark as a favorite.
- 2. Click the star icon () to add the storage object to the Favorites list.

Adding an aggregate to the Favorites list

- 1. In the left navigation pane, click **Health > Aggregates**.
- 2. In the Health/Aggregates inventory page, click the aggregate that you want to add to the Favorites list.
- 3. In the Health/Aggregate details page, click the star icon (🖈).

After you finish

To remove a storage object from the Favorites list, go to the Favorites list page, click the star icon () on the object card you want to remove, and then select the **Remove from Favorites** option.

Cluster favorite card

The Cluster favorite card enables you to view the capacity, configuration, and performance details of the individual clusters that you marked as favorites.

Cluster attributes

The Cluster favorite card displays the following attributes of individual clusters:

· Cluster health status

An icon that indicates the health of the cluster. The possible values are Normal, Warning, Error, and Critical.

· Cluster name

Name of the cluster.

Capacity

Total free space on the cluster.

Configuration

Configuration details of the cluster.

IP Address

IP address, or host name, of the cluster management logical interface (LIF) that was used to add the cluster.

Number of nodes

Number of nodes in the cluster.

Performance

Performance details of the cluster.

• IOPS

Average number of I/O operations per second over the last 72 hours.

Throughput

Average throughput over the last 72 hours, in MBps.

Aggregate favorite card

The Aggregate favorite card enables you to view the capacity and performance details of the aggregates that you marked as favorites.

Aggregate attributes

The Aggregate favorite card displays the following aggregate attributes:

Aggregate health status

An icon that indicates the health of the aggregate. The possible values are Normal, Warning, Error, and Critical.

· Aggregate name

Name of the aggregate.

Position your cursor over the aggregate name to view the name of the cluster to which the aggregate belongs.

Capacity

Percentage of free space available on the aggregate, and the estimated number of days until the aggregate becomes full.

Note that for FabricPool aggregates that this information reflects only the capacity on the local performance tier. Click the Capacity tile to view detailed information on the Health/Aggregate details page.

Performance

Performance details of the aggregate.

IOPS

Average number of I/O operations per second over the last 72 hours.

Throughput

Average throughput over the last 72 hours, in MBps.

Latency

Average response time required for an operation, in milliseconds.

Volume favorite card

The Volume favorite card enables you to view the capacity, protection, and performance details of the volumes that you marked as favorites.

Volume attributes

The Volume favorite card displays the following volume attributes:

· Volume health status

An icon that indicates the health status of the volume. The possible values are Normal, Warning, Error, and Critical.

Volume name

Name of the volume.

Capacity

Percentage of free space available on the volume, and the estimated number of days until the volume would become full.

Protection

Protection role that is set for the volume. The possible values are Unprotected, Not Applicable, Protected, and Destination.

Performance

Performance statistics for the volume.

IOPS

Average number of I/O operations per second over the last 72 hours.

Throughput

Average throughput over the last 72 hours, in MBps.

Latency

Average response time required for an operation, in milliseconds.

Creating and importing reports into Unified Manager

While Unified Manager provides reporting functionality, you might need to create new reports that are specific to your environment. You can create new reports using the Eclipse Business Intelligence and Reporting Tools (BIRT), and then import them into Unified Manager to view and manage.

Before you begin

You must have the OnCommand Administrator role.

You must have downloaded and installed MySQL Connector/J. You must have the location of the mysql-connector-java-5.1.32-bin.jar file to create the JBDC data source, which connects the report to Unified Manager.

About this task

For more detailed information on creating reports, see the Eclipse BIRT website.

Downloading and installing MySQL Connector/J

You must download and install the MySQL Connector/J drivers in a specific location. You can use these drivers to create a data source that connects the report to Unified Manager.

About this task

You must use MySQL Connector/J version 5.1 or later.

Steps

- 1. Download the MySQL Connector/J drivers at dev.mysql.com.
- Install the .jar file and note its location for future reference.

For example, install the .jar file at C:\Program Files\MySQL\MySQL Connector J\mysql-connector-java-5.1.32-bin.jar.

Creating a database user

To support a connection between Workflow Automation and Unified Manager, or to access database views, you must first create a database user with the Integration Schema or Report Schema role in the Unified Manager web UI.

Before you begin

You must have the OnCommand Administrator role.

About this task

Database users provide integration with Workflow Automation and access to report-specific database views. Database users do not have access to the Unified Manager web UI or the maintenance console, and cannot execute API calls.

Steps

- 1. In the toolbar, click , and then click Management > Users.
- 2. In the **Management/Users** page, click **Add**.
- 3. In the Add User dialog box, select Database User in the Type drop-down list.
- 4. Type a name and password for the database user.
- 5. In the **Role** drop-down list, select the appropriate role.

If you are	Choose this role
Connecting Unified Manager with Workflow Automation	Integration Schema
Accessing reporting and other database views	Report Schema

6. Click Add.

Downloading the Eclipse Business Intelligence and Reporting Tools (BIRT)

To create and import reports to Unified Manager, you must first download the Eclipse Business Intelligence and Reporting Tools (BIRT).

Steps

1. Download the BIRT software at http://download.eclipse.org/birt/downloads/.

After you finish

After downloading the BIRT software, you must extract the resulting .zip file.

Creating a project using BIRT

Before creating a report for import into Unified Manager, you must first create a project using BIRT.

Before you begin

You must have downloaded and extracted the BIRT .zip file.

Steps

- 1. From the Eclipse interface, select File > New > Project.
- 2. Expand the Business Intelligence and Reporting Tools folder, select Report Project, and click Next.
- 3. Type the project name and click **Finish**.

Creating a new report using BIRT

You can create a new report using the Eclipse plug-in for Business Intelligence and Reporting Tools (BIRT). You might want to create new reports if the existing reports in Unified Manager do not meet the needs of your environment.

Before you begin

You must have downloaded and extracted BIRT.

You must have created a project using BIRT.

Steps

- 1. From the BIRT interface, select **File > New > Report**.
- In the New Report dialog box and select the project folder, which should be the same as the project folder previously created.

If you select a different project folder, you cannot use the reporting operations in Unified Manager.

- 3. Type the report file name, and click **Next**.
- Select the report type and click Finish.

Creating a JDBC data source using BIRT

After you have created the new report using BIRT, you must create a data source to connect the report to Unified Manager.

Before you begin

You must have created a report using BIRT.

You must have downloaded and installed MySQL Connector/J.

You must have created a database user with the Report Schema role.

Steps

- 1. In Eclipse, select Data Explorer > Data Sources > New Data Source.
- 2. Select Create from a data source type in the following list.
- 3. Select JDBC Data Source, and then click Next.
- 4. In the New JDBC Data Source Profile dialog box, select com.mysql.jdbc.Driver(v5.1).
 - a. If the MySQL driver does not appear, click Manage Drivers.
 - b. In the Manage JDBC Drivers dialog box, click Add.
 - c. Browse to the location where the MySQL Connector/J .jar file was installed, and then select the file.
 - d. Click OK.

You should be able to view and select the MySQL driver.

5. Enter the fully qualified host name or the IP address of the Unified Manager instance using appropriate format:

Address Type	Format
IPv4	<pre>jdbc:mysql://xx.xx.xx.xx:3306/ocum_rep ort</pre>
IPv6	<pre>jdbc:mysql://address=(protocol=tcp) (ho st=xx:xx:xx:xx:xx:xx) (port=3306) /ocum_report</pre>

6. Enter the user name for the database user, enter the password, and then click Finish.

Creating a new MySQL data set using BIRT

After creating the data source, you must create a MySQL data set to create the output results for your report. You can also edit the output types after creating the data set.

Before you begin

You must have created a JDBC data source using BIRT.

You must have downloaded and installed MySQL Connector/J.

You must have created a database user with the Report Schema role in Unified Manager.

Steps

- 1. From **Eclipse**, select a workspace.
- Select Data Explorer > Data Sets > New Data Set.
- In the New Data Set dialog box, select the data source previously created, the data set type, and the data set name, and click Next.
- 4. Define an SQL query text using the available items, or manually enter the query, and click Finish.
- 5. Click **Preview Results** to confirm the SQL query, and then click **OK**.
- 6. In the Edit Data Set dialog box, define the output columns as necessary and click OK.
- 7. Drag items into the newly created report.

After you finish

You should now import the newly created report into Unified Manager.

Importing reports

If you have created a report outside of Unified Manager, you can import and save the report file to use with Unified Manager.

Before you begin

You must have the OnCommand Administrator role.

You must ensure that the report you plan to import is supported by Unified Manager.

Steps

- 1. In the left navigation pane, click **Reports**, and then click **Import Report**.
- In the Import Report dialog box, click Browse and select the file you want to import, and then enter a name and brief description of the report.
- 3. Click Import.

If you cannot import the report, you can check the log file to find the error causing the issue.

Using Unified Manager REST APIs

You can use REST APIs to help manage your clusters by viewing the health, capacity, and performance information captured by Unified Manager.

Accessing REST APIs using the Swagger API web page

REST APIs are exposed through the Swagger web page. You can access the Swagger web page to display the Unified Manager REST API documentation, as well as to manually issue an API call.

Before you begin

• You must have one of the following roles: Operator, Storage Administrator, or OnCommand Administrator.

You must know the IP address or fully qualified domain name of the Unified Manager server on which you
want to execute the REST APIs.

About this task

An example is provided for each REST API in the Swagger web page to help explain the objects and attributes you can use to return the information you are interested in reviewing.

Steps

1. Access the Unified Manager REST APIs.

Option	Description
From the Unified Manager web UI:	From the Menu Bar, click the Help button and then select API Documentation .
From the browser window:	Using the Unified Manager server IP address or FQDN, enter the URL to access the REST API page in the format https:// <unified_manager_ip_address_orname>/apidocs/. For example, https://10.10.10.10/apidocs/</unified_manager_ip_address_orname>

A list of API resource types, or categories, is displayed.

2. Click an API resource type to display the APIs in that resource type.

List of available REST APIs

You should be aware of the available REST APIs in Unified Manager so you can plan how you may use the APIs. The API calls are organized under the various resource types or categories.

You must refer to the Swagger web page for a complete list of the available API calls, as well as the details of each call.

The management API calls are organized according to the following categories:

- Aggregates
- Clusters
- Events
- LIFs
- LUNs
- Namespaces
- Nodes
- Ports
- SVMs

Volumes

When you select one of the categories a list appears that shows the API sub-category along with a versioned sub-category, for example:

- · /aggregates
- /v1/aggregates

The newest version of the REST APIs are listed without a version number in the URL. You should always use the newest version of the API to integrate with Unified Manager.

Setting up and monitoring an SVM with Infinite Volume without storage classes

You should use OnCommand Workflow Automation (WFA) and Unified Manager to set up and monitor storage virtual machines (SVMs) with Infinite Volume. You should create the SVM with Infinite Volume using WFA and then monitor the Infinite Volume using Unified Manager. Optionally, you can configure data protection for your Infinite Volume.

Before you begin

The following requirements must be met:

- WFA must be installed and the data sources must be configured.
- You must have the OnCommand Administrator or Storage Administrator role.
- You must have created the required number of aggregates by customizing the appropriate predefined workflow in WFA.
- You must have configured the Unified Manager server as a data source in WFA, and then you must have verified that the data is cached successfully.

About this task

- · You can monitor only data SVMs using Unified Manager.
- While performing this task, you are required to switch between two applications: OnCommand Workflow Automation (WFA) and OnCommand Unified Manager.
- The task provides high-level steps.

For details about performing the WFA tasks, see the OnCommand Workflow Automation documentation.

Steps

1. Workflow
Automation 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. 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.

3. Unified Manager Based on your organization's requirements, modify the thresholds for the Infinite Volume on the SVM.



You should use the default Infinite Volume threshold settings.

- 4. Unified Manager Configure notification alerts and traps to address any availability and capacity issues related to the Infinite Volume.
- 5. Workflow
 Automation 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.

Editing the Infinite Volume threshold settings

When you need to address any issues in your Infinite Volume's storage space, you can edit the threshold settings of the Infinite Volume's capacity based on your organization's requirements. When a threshold is crossed, events are generated, and you receive notifications if you have configured alerts for such events.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select a SVM with Infinite Volume.
- 3. In the **Health/Storage Virtual Machine** details page, click **Actions > Edit Thresholds**.
- 4. In the **Edit SVM with Infinite Volume Thresholds** dialog box, modify the thresholds as required.
- 5. Click Save and Close.

Managing your Infinite Volume with storage classes and data policies

You can effectively manage your Infinite Volume by creating the Infinite Volume with the required number of storage classes, configuring thresholds for each storage class, creating rules and a data policy to determine the placement of data written to the Infinite Volume, configuring data protection, and optionally configuring notification alerts.

Before you begin

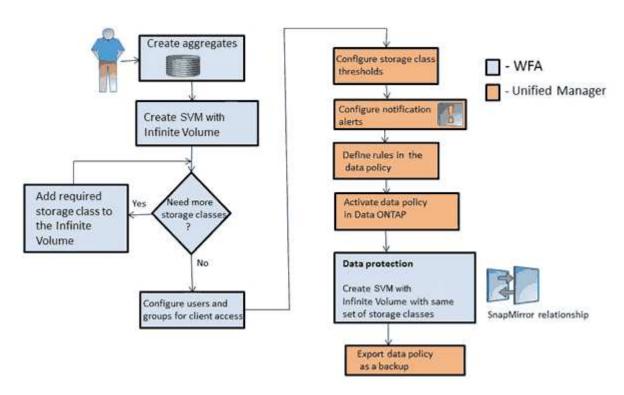
- OnCommand Workflow Automation (WFA) must be installed.
- You must have the OnCommand Administrator or Storage Administrator role.

- You must have created the required number of aggregates by customizing the appropriate predefined workflow in WFA.
- You must have created the required number of storage classes by customizing the appropriate predefined workflow in WFA.
- You must have configured the Unified Manager server as a data source in WFA, and then you must have verified that the data is cached successfully.

About this task

While performing this task, you are required to switch between two applications: OnCommand Workflow Automation (WFA) and OnCommand Unified Manager.

The task provides high-level steps. For details about performing the WFA tasks, see the *OnCommand Workflow Automation* documentation.



Steps

- Workflow
 Automation Customize the predefined workflow to define the required storage classes.
- 2. Workflow
 Automation Create an SVM with Infinite Volume with the required number of storage classes by using the appropriate workflow.
- 3. Unified Manager
 Add the cluster containing the SVM with Infinite Volume to the Unified Manager database.

You can add the cluster by providing the IP address or the FQDN of the cluster.

4. Unified Manager Based on your organization's requirements, modify the thresholds for each storage class.

You should use the default storage class threshold settings to effectively monitor storage class space.

- 5. Unified Manager Configure notification alerts and traps to address any availability and capacity issues related to the Infinite Volume.
- 6. Unified Manager Set up rules in the data policy, and then activate all the changes made to the data policy.

Rules in a data policy determine the placement of the content written to the Infinite Volume.



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. Workflow
 Automation Create a disaster recovery (DR) SVM with Infinite Volume, and then configure a data protection (DP) by performing the following steps:
 - a. Create a data protection (DP) Infinite Volume by using the appropriate workflow.
 - b. Set up a DP mirror relationship between the source and destination by using the appropriate workflow.

Editing the threshold settings of storage classes

When you need to address any issues related to storage space in your storage classes, you can edit the threshold settings of the storage class capacity based on your organization's requirements. When the threshold is crossed, events are generated, and you receive notifications if you have configured alerts for such events.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- 2. In the Health/Storage Virtual Machines inventory page, select an SVM with Infinite Volume.
- 3. In the **Health/Storage Virtual Machine** details page, click **Actions > Edit Thresholds**.
- 4. In the **Edit Storage Class Thresholds** dialog box, modify the thresholds as required.
- 5. Click Save and Close.

Adding alerts

You can configure alerts to notify you when a particular event is generated. You can configure alerts for a single resource, for a group of resources, or for events of a particular severity type. You can specify the frequency with which you want to be notified and associate a script to the alert.

Before you begin

 You must have configured notification settings such as the user email address, SMTP server, and SNMP trap host to enable the Unified Manager server to use these settings to send notifications to users when an event is generated.

- You must know the resources and events for which you want to trigger the alert, and the user names or email addresses of the users that you want to notify.
- If you want to have a script execute based on the event, you must have added the script to Unified Manager by using the Management/Scripts page.
- You must have the OnCommand Administrator or Storage Administrator role.

About this task

You can create an alert directly from the Event details page after receiving an event in addition to creating an alert from the Configuration/Alerting page, as described here.

Steps

- 1. In the left navigation pane, click **Configuration > Alerting**.
- 2. In the **Configuration/Alerting** page, click **Add**.
- 3. In the Add Alert dialog box, click Name, and enter a name and description for the alert.
- 4. Click **Resources**, and select the resources to be included in or excluded from the alert.

You can set a filter by specifying a text string in the **Name contains** field to select a group of resources. Based on the text string that you specify, the list of available resources displays only those resources that match the filter rule. The text string that you specify is case-sensitive.

If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.

5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.



To select more than one event, press the Ctrl key while you make your selections.

Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.



If you modify the email address that is specified for the user and reopen the alert for editing, the Name field appears blank because the modified email address is no longer mapped to the user that was previously selected. Also, if you modified the email address of the selected user from the Management/Users page, the modified email address is not updated for the selected user.

You can also choose to notify users through SNMP traps.

7. Click Save.

Example of adding an alert

This example shows how to create an alert that meets the following requirements:

· Alert name: HealthTest

- Resources: includes all volumes whose name contains "abc" and excludes all volumes whose name contains "xyz"
- Events: includes all critical health events
- Actions: includes "sample@domain.com", a "Test" script, and the user has to be notified every 15 minutes

Perform the following steps in the Add Alert dialog box:

- 1. Click Name, and enter HealthTest in the Alert Name field.
- 2. Click Resources, and in the Include tab, select Volumes from the drop-down list.
 - a. Enter abc in the Name contains field to display the volumes whose name contains "abc".
 - b. Select << All Volumes whose name contains 'abc'>> from the Available Resources area, and move it to the Selected Resources area.
 - c. Click **Exclude**, and enter xyz in the **Name contains** field, and then click **Add**.
- 3. Click **Events**, and select **Critical** from the Event Severity field.
- 4. Select All Critical Events from the Matching Events area, and move it to the Selected Events area.
- 5. Click Actions, and enter sample@domain.com in the Alert these users field.
- 6. Select **Remind every 15 minutes** to notify the user every 15 minutes.

You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.

- 7. In the Select Script to Execute menu, select **Test** script .
- 8. Click Save.

Creating rules

You can add new rules to your data policy to determine the placement of data that is written to the Infinite Volume. You can create rules either by using rule templates that are defined in Unified Manager or create custom rules.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

Creating rules using templates

You can add new rules by using rule templates defined by Unified Manager to determine the placement of data that is written to the SVM with Infinite Volume. You can create rules based on file types, directory paths, or owners.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- 3. Click the Data Policy tab.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

- 4. Click Create.
- 5. In the Create Rule dialog box, choose an appropriate rule template from the drop-down list.

The template is based on three categories: file type, owner, or directory path.

- 6. Based on the template selected, add the necessary conditions in the Matching Criteria area.
- 7. Select an appropriate storage class from the **Place the matching content in Storage Class** drop-down list.
- 8. Click Create.

The new rule you created is displayed in the Data Policy tab.

- 9. Preview any other changes made to the data policy.
- 10. Click **Activate** to activate the changes in the rule properties in the SVM.

Creating custom rules

Based on your data center requirements, you can create custom rules and add them to a data policy to determine the placement of data that is written to the SVM with Infinite Volume. You can create custom rules from the Create Rule dialog box without using any existing template.

Before you begin

- You must have the OnCommand Administrator or Storage Administrator role.
- The cluster containing the SVM with Infinite Volume with storage classes must be added to the Unified Manager database.

About this task

The Data Policy tab is visible only for an SVM with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- 3. Click Data Policy.

- 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. Preview any other changes made to the data policy.
- 10. Click **Activate** to activate the changes in the rule properties in the SVM.

Exporting a data policy configuration

You can export a data policy configuration from Unified Manager to a file. For example, after you have taken the required backup, and in the event of a disaster, you can export the data policy configuration from the primary.

Before you begin

You must have the OnCommand Administrator or Storage Administrator role.

About this task

The Data Policy tab, which is used while performing this task, is displayed only for SVMs with Infinite Volume.

Steps

- 1. In the left navigation pane, click **Health > SVMs**.
- In the Health/Storage Virtual Machines inventory page, select the appropriate SVM.
- Click Data Policy.

The list of rules in the data policy for the selected SVM with Infinite Volume is displayed.

- 4. Click **Export**.
- In the browser-specific dialog box, specify the location to which the data policy configuration has to be exported.

Results

The data policy configuration is exported as a JSON file in the specified location.

Sending a Unified Manager support bundle to technical support

This workflow shows you how to generate, retrieve, and send a support bundle to technical support using the Unified Manager maintenance console. You should send a

support bundle when the issue that you have requires more detailed diagnosis and troubleshooting than an AutoSupport message provides.

About this task

For more information about the maintenance console and support bundles, see Using the maintenance console.

Unified Manager stores two generated support bundles at one time.

Accessing the maintenance console

If the Unified Manager user interface is not in operation, or if you need to perform functions that are not available in the user interface, you can access the maintenance console to manage your Unified Manager system.

Before you begin

You must have installed and configured Unified Manager.

About this task

After 15 minutes of inactivity, the maintenance console logs you out.



When installed on VMware, if you have already logged in as the maintenance user through the VMware console, you cannot simultaneously log in using Secure Shell.

Steps

1. Follow these steps to access the maintenance console:

On this operating system	Follow these steps
VMware	 Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance.
	b. Log in to the maintenance console using your maintenance user name and password.
Linux	Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system.
	b. Log in to the system with the maintenance user (umadmin) name and password.
	c. Enter the command maintenance_console and press Enter.

On this operating system	Follow these steps
Windows	a. Log in to the Unified Manager system with administrator credentials.
	b. Launch PowerShell as a Windows administrato
	c. Enter the command maintenance_console and press Enter.
	On Microsoft Windows Server 2012 if you receive an execution policy error, enter the following command and try step c again: PowerShell.exe -ExecutionPolicy RemoteSigned

The Unified Manager maintenance console menu is displayed.

Generating a support bundle

You can generate a support bundle, containing full diagnostic information, so that you can then retrieve it and send it to technical support for troubleshooting help. Because some types of data can use a large amount of cluster resources or take a long time to complete, you can specify data types to include or exclude in the support bundle.

Before you begin

You must have access to the maintenance console as the maintenance user.

About this task

Unified Manager stores only the two most recently generated support bundles. Older support bundles are deleted from the system.



On Windows systems, the command supportbundle.bat is no longer supported to generate a support bundle.

Steps

- 1. In the maintenance console **Main Menu**, select **Support/Diagnostics**.
- 2. Select Generate Support Bundle.
- 3. Select or deselect the following data types to include or exclude in the support bundle:
 - database dump

A dump of the MySQL Server database.

heap dump

A snapshot of the state of the main Unified Manager server processes. This option is disabled by default and should be selected only when requested by customer support.

acquisition recordings

A recording of all communications between Unified Manager and the monitored clusters.



If you deselect all data types, the support bundle is still generated with other Unified Manager data.

4. Type g, and then press Enter to generate the support bundle.

Since the generation of a support bundle is a memory intensive operation, you are prompted to verify that you are sure you want to generate the support bundle at this time.

Type y, and then press Enter to generate the support bundle.

If you do not want to generate the support bundle at this time, type n, and then press Enter.

- 6. If you included database dump files in the support bundle, you are prompted to specify the time period for which you want performance statistics included. Including performance statistics can take a lot of time and space, so you can also dump the database without including performance statistics:
 - a. Enter the starting date in the format YYYYMMDD.

For example, enter 20170101 for January 1, 2017. Enter n if you do not want performance statistics to be included.

b. Enter the number of days of statistics to include, beginning from 12 a.m. on the specified starting date.

You can enter a number from 1 through 10.

If you are including performance statistics, the system displays the period of time for which performance statistics will be collected.

7. Select Generate Support Bundle.

The generated support bundle resides in the /support directory.

After you finish

After generating the support bundle, you can retrieve it using an SFTP client or by using UNIX or Linux CLI commands. On Windows installations you can use Remote Desktop (RDP) to retrieve the support bundle.

The generated support bundle resides in the /support directory on VMware systems, in /opt/netapp/data/support/ on Linux systems, and in ProgramData\NetApp\OnCommandAppData\ocum\support on Windows systems.

Retrieving the support bundle using a Windows client

If you are a Windows user, you can download and install a tool to retrieve the support bundle from your Unified Manager server. You can send the support bundle to technical support for a more detailed diagnosis of an issue. Filezilla or WinSCP are examples of tools you can use.

Before you begin

You must be the maintenance user to perform this task.

You must use a tool that supports SCP or SFTP.

Steps

- 1. Download and install a tool to retrieve the support bundle.
- 2. Open the tool.
- 3. Connect to your Unified Manager management server over SFTP.

The tool displays the contents of the /support directory and you can view all existing support bundles.

- 4. Select the destination directory for the support bundle you want to copy.
- 5. Select the support bundle you want to copy and use the tool to copy the file from the Unified Manager server to your local system.

Related information

Filezilla - https://filezilla-project.org/

WinSCP - http://winscp.net

Retrieving the support bundle using a UNIX or Linux client

If you are a UNIX or Linux user, you can retrieve the support bundle from your vApp by using the command-line interface (CLI) on your Linux client server. You can use either SCP or SFTP to retrieve the support bundle.

Before you begin

You must be the maintenance user to perform this task.

You must have generated a support bundle using the maintenance console and have the support bundle name available.

Steps

- 1. Access the CLI through Telnet or the console, using your Linux client server.
- 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	<pre>scp <maintenance-user>@<vapp-name-or- ip="">:/support/support_bundle_file_name. 7z <destination-directory></destination-directory></vapp-name-or-></maintenance-user></pre>

If you are using	Then use the following command
SFTP	<pre>sftp <maintenance-user>@<vapp-name-or- ip="">:/support/support_bundle_file_name. 7z <destination-directory></destination-directory></vapp-name-or-></maintenance-user></pre>

The name of the support bundle is provided to you when you generate it using the maintenance console.

4. Enter the maintenance user password.

Examples

The following example uses SCP to retrieve the support bundle:

```
$ scp admin@10.10.12.69:/support/support_bundle_20160216_145359.7z
.
Password: <maintenance_user_password>
support_bundle_20160216_145359.7z 100% 119MB 11.9MB/s 00:10
```

The following example uses SFTP to retrieve the support bundle:

```
$ sftp
admin@10.10.12.69:/support/support_bundle_20160216_145359.7z .
Password: <maintenance_user_password>
Connected to 10.228.212.69.
Fetching /support/support_bundle_20130216_145359.7z to
./support_bundle_20130216_145359.7z
/support/support_bundle_20160216_145359.7z
```

Sending a support bundle to technical support

When an issue requires more detailed diagnosis and troubleshooting information than an AutoSupport message provides, you can send a support bundle to technical support.

Before you begin

You must have access to the support bundle to send it to technical support.

You must have a case number generated through the technical support web site.

Steps

- 1. Log in to the NetApp Support Site.
- 2. Upload the file.

How to upload a file to NetApp

Tasks and information related to several workflows

Some tasks and reference texts that can help you understand and complete a workflow are common to many of the workflows in Unified Manager, including adding and reviewing notes about an event, assigning an event, acknowledging and resolving events, and details about volumes, storage virtual machines (SVMs), aggregates, and so on.

Adding and reviewing notes about an event

While addressing events, you can add information about how the issue is being addressed by using the Notes and Updates area in the Event details page. This information can enable another user who is assigned to address the event. You can also view information that was added by the user who last addressed an event, based on the recent timestamp.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Events**.
- 2. From the **Events** inventory page, click the event for which you want to add the event-related information.
- 3. In the **Event** details page, add the required information in the **Notes and Updates** area.
- 4. Click Post.

Assigning events to specific users

You can assign unassigned events to yourself or to other users, including remote users. You can reassign assigned events to another user, if required. For example, when frequent issues occur on a storage object, you can assign the events for these issues to the user who manages that object.

Before you begin

- The user's name and email ID must be configured correctly.
- You must have the Operator, OnCommand Administrator, or Storage Administrator role.

Steps

- 1. In the left navigation pane, click **Events**.
- In the Events inventory page, select one or more events that you want to assign.
- 3. Assign the event by choosing one of the following options:

If you want to assign the event to	Then do this
Yourself	Click Assign To > Me.

Then do this
a. Click Assign To > Another user.
b. In the Assign Owner dialog box, enter the user name, or select a user from the drop-down list.
c. Click Assign .
An email notification is sent to the user.
If you do not enter a user name or select a user from the dropdown list, and click Assign , the event remains unassigned.

Acknowledging and resolving events

You should acknowledge an event before you start working on the issue that generated the event so that you do not continue to receive repeat alert notifications. After you take corrective action for a particular event, you should mark the event as resolved.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

You can acknowledge and resolve multiple events simultaneously.



You cannot acknowledge Information events.

Steps

- 1. In the left navigation pane, click **Events**.
- 2. From the events list, perform the following actions to acknowledge the events:

If you want to	Do this
Acknowledge and mark a single event as resolved	a. Click the event name.b. From the Event details page, determine the cause of the event.c. Click Acknowledge.
	d. Take appropriate corrective action. e. Click Mark As Resolved.

If you want to	Do this
Acknowledge and mark multiple events as resolved	a. Determine the cause of the events from the respective Event details page.
	b. Select the events.
	c. Click Acknowledge .
	d. Take appropriate corrective actions.
	e. Click Mark As Resolved.

After the event is marked resolved, the event is moved to the resolved events list.

3. 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 possible remediations to resolve the issue.

Event Name

The name of the event and the time the event was last seen.

For non-performance events, while the event is in the New or Acknowledged state the last seen information is not known and is therefore hidden.

Event Description

A brief description of the event.

In some cases a reason for the event being triggered is provided in the event description.

Component in Contention

For dynamic performance events, this section displays icons that represent the logical and physical components of the cluster. If a component is in contention, its icon is circled and highlighted red.

The following components may be displayed:

Network

Represents the wait time of I/O requests by the iSCSI protocols or the Fibre Channel (FC) protocols on the cluster. The wait time is time spent waiting for iSCSI Ready to Transfer (R2T) or FCP Transfer Ready (XFER_RDY) transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the block protocol layer is impacting the latency of one or more workloads.

Network Processing

Represents the software component in the cluster involved with I/O processing between the protocol layer and the cluster. The node handling network processing might have changed since the event was detected. If the network processing component is in contention, it means high utilization at the network

processing node is impacting the latency of one or more workloads.

QoS Policy

Represents the storage Quality of Service (QoS) policy group of which the workload is a member. If the policy group component is in contention, it means all workloads in the policy group are being throttled by the set throughput limit, which is impacting the latency of one or more of those workloads.

Cluster Interconnect

Represents the cables and adapters with which clustered nodes are physically connected. If the cluster interconnect component is in contention, it means high wait time for I/O requests at the cluster interconnect is impacting the latency of one or more workloads.

Data Processing

Represents the software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains the workload. The node handling data processing might have changed since the event was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the latency of one or more workloads.

MetroCluster Resources

Represents the MetroCluster resources, including NVRAM and interswitch links (ISLs), used to mirror data between clusters in a MetroCluster configuration. If the MetroCluster component is in contention, it means high write throughput from workloads on the local cluster or a link health issue is impacting the latency of one or more workloads on the local cluster. If the cluster is not in a MetroCluster configuration, this icon is not displayed.

Aggregate or SSD Aggregate Ops

Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate). An "SSD Aggregate" consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

Cloud Latency

Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

Sync SnapMirror

Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

The Event Information, System Diagnosis, and Suggested Actions sections are described in other topics.

Command buttons

The command buttons enable you to perform the following tasks:

Notes icon

Enables you to add or update a note about the event, and review all notes left by other users.

Actions menu

· Assign to Me

Assigns the event to you.

Assign to Others

Opens the Assign Owner dialog box, which enables you to assign or reassign the event to other users.

When you assign an event to a user, the user's name and the time when the event was assigned are added in the events list for the selected events.

You can also unassign events by leaving the ownership field blank.

Acknowledge

Acknowledges the selected events so that you do not continue to receive repeat alert notifications.

When you acknowledge an event, your user name and the time that you acknowledged the event are added in the events list (Acknowledged By) for the selected events. When you acknowledge an event, you take responsibility for managing that event.

Mark As Resolved

Enables you to change the event state to Resolved.

When you resolve an event, your user name and the time that you resolved the event are added in the events list (Resolved By) for the selected events. After you have taken corrective action for the event, you must mark the event as resolved.

Add Alert

Displays the Add Alert dialog box, which enables you to add an alert for the selected event.

What the Event Information section displays

You use the Event Information section on the Event details page to view the details about a selected event, such as the event severity, impact level, impact area, and event source.

Fields that are not applicable to the event type are hidden. You can view the following event details:

Event Trigger Time

The time at which the event was generated.

State

The event state: New, Acknowledged, Resolved, or Obsolete.

Obsoleted Cause

The actions that caused the event to be obsoleted, for example, the issue was fixed.

Event Duration

For active (new and acknowledged) events, this is the time between detection and the time when the event was last analyzed. For obsolete events, this is the time between detection and when the event was resolved.

This field is displayed for all performance events, and for other event types only after they have been resolved or obsoleted.

Last Seen

The date and time at which the event was last seen as active.

For performance events this value may be more recent than the Event Trigger Time as this field is updated after each new collection of performance data as long as the event is active. For other types of events, when in the New or Acknowledged state, this content is not updated and the field is therefore hidden.

Severity

The event severity: Critical (\mathbf{X}), Error ($\mathbf{0}$), Warning ($\mathbf{\Lambda}$), and Information ($\mathbf{0}$).

Impact Level

The event impact level: Incident, Risk, or Event.

Impact Area

The event impact area: Availability, Capacity, Performance, Protection, or Configuration.

Source

The name of the object on which the event has occurred.

When viewing the details for a shared QoS policy event, up to three of the workload objects that are consuming the most IOPS or MBps are listed in this field.

You can click the source name link to display the health or performance details page for that object.

Source Annotations

Displays the annotation name and value for the object to which the event is associated.

This field is displayed only for health events on clusters, SVMs, and volumes.

Source Groups

Displays the names of all the groups of which the impacted object is a member.

This field is displayed only for health events on clusters, SVMs, and volumes.

Source Type

The object type (for example, SVM, Volume, or Qtree) with which the event is associated.

On Cluster

The name of the cluster on which the event occurred.

You can click the cluster name link to display the health or performance details page for that cluster.

Affected Objects Count

The number of objects affected by the event.

You can click the object link to display the inventory page populated with the objects that are currently affected by this event.

This field is displayed only for performance events.

Affected Volumes

The number of volumes that are being affected by this event.

This field is displayed only for performance events on nodes or aggregates.

Triggered Policy

The name of the threshold policy that issued the event.

You can hover your cursor over the policy name to see the details of the threshold policy. For adaptive QoS policies the defined policy, block size, and allocation type (allocated space or used space) is also displayed.

This field is displayed only for performance events.

Acknowledged by

The name of the person who acknowledged the event and the time that the event was acknowledged.

Resolved by

The name of the person who resolved the event and the time that the event was resolved.

Assigned to

The name of the person who is assigned to work on the event.

Alert Settings

The following information about alerts is displayed:

• If there are no alerts associated with the selected event, an Add alert 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 Configuration/Alerting page by clicking the link to view more details about these alerts.

Alerts that are disabled are not displayed.

Last Notification Sent

The date and time at which the most recent alert notification was sent.

Sent Via

The mechanism that was used to send the alert notification: email or SNMP trap.

Previous Script Execution

The name of the script that was executed when the alert was generated.

What the System Diagnosis section displays

The System Diagnosis section of the Event details page provides information that can help you diagnose issues that may have been responsible for the event.

This area is displayed only for some events.

Some performance events provide charts that are relevant to the particular event that has been triggered. Typically this includes and IOPS or MBps chart and a latency chart for the previous ten days. When arranged this way you can see which storage components are most affecting latency, or being affected by latency, when the event is active.

For dynamic performance events, the following charts are displayed:

- Workload Latency Displays the history of latency for the top victim, bully, or shark workloads at the component in contention.
- Workload Activity Displays details about the workload usage of the cluster component in contention.
- Resource Activity Display historical performance statistics for the cluster component in contention.

Other charts are displayed when some cluster components are in contention.

Other events provide a brief description of the type of analysis the system is performing on the storage object. In some cases there will be one or more lines; one for each component that has been analyzed, for system-defined performance policies that analyze multiple performance counters. In this scenario, a green or red icon displays next to the diagnosis to indicate whether an issue was found, or not, in that particular diagnosis.

What the Suggested Actions section displays

The Suggested Actions section of the Event details page provides possible reasons for the event and suggests a few actions so that you can try to resolve the event on your own. The suggested actions are customized based on the type of event or type of threshold that has been breached. This area is displayed only for some types of events.

In some cases there are **Help** links provided on the page that reference additional information for many suggested actions, including instructions for performing a specific action. Some of the actions may involve using Unified Manager, OnCommand System Manager, OnCommand Workflow Automation, ONTAP CLI commands, or a combination of these tools.

There are also some links provided in this help topic.

You should consider the actions suggested here as only a guidance in resolving this event. The action you take to resolve this event should be based on the context of your environment.

Description of event severity types

Each event is associated with a severity type to help you prioritize the events that require immediate corrective action.

Critical

A problem occurred that might lead to service disruption if corrective action is not taken immediately.

Performance critical events are sent from user-defined thresholds only.

Error

The event source is still performing; however, corrective action is required to avoid service disruption.

Warning

The event source experienced an occurrence that you should be aware of, or a performance counter for a cluster object is out of normal range and should be monitored to make sure it does not reach the critical severity. Events of this severity do not cause service disruption, and immediate corrective action might not be required.

Performance warning events are sent from user-defined, system-defined, or dynamic thresholds.

Information

The event occurs when a new object is discovered, or when a user action is performed. For example, when any storage object is deleted or when there are any configuration changes, the event with severity type Information is generated.

Information events are sent directly from ONTAP when it detects a configuration change.

Description of event impact levels

Each event is associated with an impact level (Incident, Risk, or Event) to help you prioritize the events that require immediate corrective action.

Incident

An incident is a set of events that can cause a cluster to stop serving data to the client and run out of space for storing data. Events with an impact level of Incident are the most severe. Immediate corrective action should be taken to avoid service disruption.

Risk

A risk is a set of events that can potentially cause a cluster to stop serving data to the client and run out of space for storing data. Events with an impact level of Risk can cause service disruption. Corrective action might be required.

Event

An event is a state or status change of storage objects and their attributes. Events with an impact level of Event are informational and do not require corrective action.

Description of event impact areas

Events are categorized into five impact areas (availability, capacity, configuration, performance, and protection) to enable you to concentrate on the types of events for which you are responsible.

Availability

Availability events notify you if a storage object goes offline, if a protocol service goes down, if an issue with storage failover occurs, or if an issue with hardware occurs.

Capacity

Capacity events notify you if your aggregates, volumes, LUNs, or namespaces are approaching or have reached a size threshold, or if the rate of growth is unusual for your environment.

Configuration

Configuration events inform you of the discovery, deletion, addition, removal, or renaming of your storage objects. Configuration events have an impact level of Event and a severity type of Information.

Performance

Performance events notify you of resource, configuration, or activity conditions on your cluster that might adversely affect the speed of data storage input or retrieval on your monitored storage objects.

Protection

Protection events notify you of incidents or risks involving SnapMirror relationships, issues with destination capacity, problems with SnapVault relationships, or issues with protection jobs. Any ONTAP object (especially aggregates, volumes, and SVMs) that host secondary volumes and protection relationships are categorized in the protection impact area.

Health/Volume details page

You can use the Health/Volume details page to view detailed information about a selected volume, such as capacity, storage efficiency, configuration, protection, annotation, and events generated. You can also view information about the related objects and related alerts for that volume.

You must have the OnCommand Administrator or Storage Administrator role.

Command buttons

The command buttons enable you to perform the following tasks for the selected volume:

Switch to Performance View

Enables you to navigate to the Performance/Volume details page.



Enables you to add the selected volume to the Favorites dashboard.

Actions

Add Alert

Enables you to add an alert to the selected volume.

· Edit Thresholds

Enables you to modify the threshold settings for the selected volume.

Annotate

Enables you to annotate the selected volume.

Protect

Enables you to create either SnapMirror or SnapVault relationships for the selected volume.

Relationship

Enables you to execute the following protection relationship operations:

Edit

Launches the Edit Relationship dialog box which enables you to change existing SnapMirror policies, schedules, and maximum transfer rates for an existing protection relationship.

Abort

Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer.

Quiesce

Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress must complete before the relationship is quiesced.

Break

Breaks the relationship between the source and destination volumes and changes the destination to a read-write volume.

Remove

Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed. This operation cannot be undone.

Resume

Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used, if one exists.

Resynchronize

Enables you to resynchronize a previously broken relationship.

Initialize/Update

Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform a manual update if the relationship is already initialized.

Reverse Resync

Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by the contents on the destination, and any data that is newer than the data on the common Snapshot copy is deleted.

Restore

Enables you to restore data from one volume to another volume.



The Restore button and the Relationship operation buttons are not available for FlexGroup volumes, or for volumes that are in synchronous protection relationships.

View Volumes

Enables you to navigate to the Health/Volumes inventory page.

Capacity tab

The Capacity tab displays details about the selected volume, such as its physical capacity, logical capacity, threshold settings, quota capacity, and information about any volume move operation:

Capacity Physical

Details the physical capacity of the volume:

Snapshot Overflow

Displays the data space that is consumed by the Snapshot copies.

Used

Displays the space used by data in the volume.

Warning

Indicates that the space in the volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

• Error

Indicates that the space in the volume is full. If this threshold is breached, the Space Full event is generated.

Unusable

Indicates that the Thin-Provisioned Volume Space At Risk event is generated and that the space in the thinly provisioned volume is at risk because of aggregate capacity issues. The unusable capacity is displayed only for thinly provisioned volumes.

· Data graph

Displays the total data capacity and the used data capacity of the volume.

If autogrow is enabled, the data graph also displays the space available in the aggregate. The data graph displays the effective storage space that can be used by data in the volume, which can be one of the following:

- Actual data capacity of the volume for the following conditions:
 - Autogrow is disabled.
 - Autogrow-enabled volume has reached the maximum size.
 - Autogrow-enabled thickly provisioned volume cannot grow further.
- Data capacity of the volume after considering the maximum volume size (for thinly provisioned volumes and for thickly provisioned volumes when the aggregate has space for the volume to reach maximum size)
- Data capacity of the volume after considering the next possible autogrow size (for thickly provisioned volumes that have an autogrow percentage threshold)
- 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.

Capacity Logical

Displays the logical space characteristics of the volume. The logical space indicates the real size of the data that is being stored on disk without applying the savings from using ONTAP storage efficiency technologies.

Logical Space Reporting

Displays if the volume has logical space reporting configured. The value can be Enabled, Disabled, or Not applicable. "Not applicable" is displayed for volumes on older versions of ONTAP or on volumes that do not support logical space reporting.

Used

Displays the amount of logical space that is being used by data in the volume, and the percentage of

logical space used based on the total data capacity.

· Available

Displays the amount of logical space that is still available for data in the volume, and the percentage of logical space available based on the total data capacity.

Logical Space Enforcement

Displays whether logical space enforcement is configured for thinly provisioned volumes. When set to Enabled, the logical used size of the volume cannot be greater than the currently set physical volume size.

Autogrow

Displays whether the volume automatically grows when it is out of space.

Space Guarantee

Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate. These blocks are then guaranteed to be available for writes to files in the volume. The space guarantee can be set to one of the following:

None

No space guarantee is configured for the volume.

File

Full size of sparsely written files (for example, LUNs) is guaranteed.

Volume

Full size of the volume is guaranteed.

Partial

The FlexCache volume reserves space based on its size. If the FlexCache volume's size is 100 MB or more, the minimum space guarantee is set to 100 MB by default. If the FlexCache volume's size is less than 100 MB, the minimum space guarantee is set to the FlexCache volume's size. If the FlexCache volume's size is grown later, the minimum space guarantee is not incremented.



The space guarantee is Partial when the volume is of type Data-Cache.

· Details (Physical)

Displays the physical characteristics of the volume.

Total Capacity

Displays the total physical capacity in the volume.

Data Capacity

Displays the amount of physical space used by the volume (used capacity) and the amount of physical space that is still available (free capacity) in the volume. These values are also displayed as a percentage

of the total physical capacity.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

Snapshot Reserve

Displays the amount of space used by the Snapshot copies (used capacity) and amount of space available for Snapshot copies (free capacity) in the volume. These values are also displayed as a percentage of the total snapshot reserve.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the Snapshot copies (used capacity) and the amount of space that is available in the volume but cannot be used for making Snapshot copies (unusable capacity) because of aggregate capacity issues is displayed.

Volume Thresholds

Displays the following volume capacity thresholds:

Nearly Full Threshold

Specifies the percentage at which a volume is nearly full.

Full Threshold

Specifies the percentage at which a volume is full.

Other Details

Autogrow Max Size

Displays the maximum size up to which the 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.

Qtree Quota Committed Capacity

Displays the space reserved in the quotas.

Qtree Quota Overcommitted Capacity

Displays the amount of space that can be used before the system generates the Volume Qtree Quota Overcommitted event.

Fractional Reserve

Controls the size of the overwrite reserve. By default, the fractional reserve is set to 100, indicating that 100 percent of the required reserved space is reserved so that the objects are fully protected for overwrites. If the fractional reserve is less than 100 percent, the reserved space for all the space-reserved files in that volume is reduced to the fractional reserve percentage.

Snapshot Daily Growth Rate

Displays the change (in percentage, or in KB, MB, GB, and so on) that occurs every 24 hours in the Snapshot copies in the selected volume.

Snapshot Days to Full

Displays the estimated number of days remaining before the space reserved for the Snapshot copies in the volume reaches the specified threshold.

The Snapshot Days to Full field displays a Not Applicable value when the growth rate of the Snapshot copies in the volume is zero or negative, or when there is insufficient data to calculate the growth rate.

Snapshot Autodelete

Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.

Snapshot Copies

Displays information about the Snapshot copies in the volume.

The number of Snapshot copies in the volume is displayed as a link. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

The Snapshot copy count is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

Volume Move

Displays the status of either the current or the last volume move operation that was performed on the volume, and other details, such as the current phase of the volume move operation which is in progress, source aggregate, destination aggregate, start time, end time, and estimated end time.

Also displays the number of volume move operations that are performed on the selected volume. You can view more information about the volume move operations by clicking the **Volume Move History** link.

Efficiency tab

The Efficiency tab displays information about the space saved in the volumes by using storage efficiency features such as deduplication, compression, and FlexClone volumes.

Deduplication

Enabled

Specifies whether deduplication is enabled or disabled on a volume.

Space Savings

Displays the amount of space saved (in percentage, or in KB, MB, GB, and so on) in a volume by using deduplication.

Last Run

Displays the time that has elapsed since the deduplication operation was last performed. Also specifies whether the deduplication operation was successful.

If the time elapsed exceeds a week, the timestamp representing when the operation was performed is displayed.

Mode

Specifies whether the deduplication operation enabled on a volume is a manual, scheduled, or policy-based operation. If the mode is set to Scheduled, the operation schedule is displayed, and if the mode is set to a policy, the policy name is displayed.

Status

Displays the current status of the deduplication operation. The status can be Idle, Initializing, Active, Undoing, Pending, Downgrading, or Disabled.

Type

Specifies the type of deduplication operation running on the volume. If the volume is in a SnapVault relationship, the type displayed is SnapVault. For any other volume, the type is displayed as Regular.

Compression

Enabled

Specifies whether compression is enabled or disabled on a volume.

Space Savings

Displays the amount of space saved (in percentage, or in KB, MB, GB, and so on) in a volume by using compression.

Configuration tab

The Configuration tab displays details about the selected volume, such as the export policy, RAID type, capacity and storage efficiency related features of the volume:

Overview

Full Name

Displays the full name of the volume.

Aggregates

Displays the name of the aggregate on which the volume resides, or the number of aggregates on which the FlexGroup volume resides.

Tiering Policy

Displays the tiering policy set for the volume; if the volume is deployed on a FabricPool-enabled aggregate. The policy can be None, Snapshot Only, Backup, or Auto.

Storage Virtual Machine

Displays the name of the storage virtual machine (SVM) that contains the volume.

Junction Path

Displays the status of the path, which can be active or inactive. The path in the SVM to which the volume is mounted is also displayed. You can click the **History** link to view the most recent five

changes to the junction path.

Export policy

Displays the name of the export policy that is created for the volume. You can click the link to view details about the export policies, authentication protocols, and access enabled on the volumes that belong to the SVM.

· Style

Displays the volume style. The volume style can be FlexVol or FlexGroup.

Type

Displays the type of the selected volume. The volume type can be Read-write, Load-sharing, Data-Protection, Data-cache, or Temporary.

RAID Type

Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, or RAID-TEC.



Multiple RAID types may display for FlexGroup volumes because the constituent volumes for FlexGroups can be on aggregates of different types.

SnapLock Type

Displays the SnapLock Type of the aggregate that contains the volume.

SnapLock Expiry

Displays the expiry date of SnapLock volume.

Capacity

Thin Provisioning

Displays whether thin provisioning is configured for the volume.

Autogrow

Displays whether the flexible volume grows automatically within an aggregate.

Snapshot Autodelete

Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.

Quotas

Specifies whether the quotas are enabled for the volume.

Efficiency

Deduplication

Specifies whether deduplication is enabled or disabled for the selected volume.

Compression

Specifies whether compression is enabled or disabled for the selected volume.

Protection

Snapshot Copies

Specifies whether automatic Snapshot copies are enabled or disabled.

Protection tab

The Protection tab displays protection details about the selected volume, such as lag information, relationship type, and topology of the relationship.

Summary

Displays SnapMirror and SnapVault relationships properties for a selected volume. For any other relationship type, only the Relationship Type property is displayed. If a primary volume is selected, only the Managed and Local Snapshot copy Policy are displayed. Properties displayed for SnapMirror and SnapVault relationships include the following:

Source Volume

Displays the name of the selected volume's source if the selected volume is a destination.

Lag Status

Displays the update or transfer lag status for a protection relationship. The status can be Error, Warning, or Critical.

The lag status is not applicable for synchronous relationships.

Lag Duration

Displays the time by which the data on the mirror lags behind the source.

Last Successful Update

Displays the date and time of the most recent successful protection update.

The last successful update is not applicable for synchronous relationships.

Storage Service Member

Displays either Yes or No to indicate whether or not the volume belongs to and is managed by a storage service.

Version Flexible Replication

Displays either Yes, Yes with backup option, or None. Yes indicates that SnapMirror replication is possible even if source and destination volumes are running different versions of ONTAP software. Yes with backup option indicates the implementation of SnapMirror protection with the ability to retain multiple versions of backup copies on the destination. None indicates that Version Flexible Replication is not enabled.

Relationship Capability

Indicates the ONTAP capabilities available to the protection relationship.

Protection Service

Displays the name of the protection service if the relationship is managed by a protection partner application.

Relationship Type

Displays any relationship type, including Asynchronous Mirror, Asynchronous Vault, StrictSync, and Sync.

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:

Aborting

SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

Checking

The destination volume is undergoing a diagnostic check and no transfer is in progress.

Finalizing

SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

Idle

Transfers are enabled and no transfer is in progress.

In-Sync

The data in the two volumes in the synchronous relationship are synchronized.

Out-of-Sync

The data in the destination volume is not synchronized with the source volume.

Preparing

SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.

Queued

SnapMirror transfers are enabled. No transfers are in progress.

Quiesced

SnapMirror transfers are disabled. No transfer is in progress.

Quiescing

A SnapMirror transfer is in progress. Additional transfers are disabled.

Transferring

SnapMirror transfers are enabled and a transfer is in progress.

Transitioning

The asynchronous transfer of data from the source to the destination volume is complete, and the transition to synchronous operation has started.

Waiting

A SnapMirror transfer has been initiated, but some associated tasks are waiting to be queued.

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 Asynchronous Mirror protection policy, and XDPDefault indicates the default Asynchronous Vault policy. StrictSync indicates the default Synchronous Strict protection policy, and Sync indicates the default Synchronous 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. 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.

Double bold lines specify an Asynchronous Mirror relationship, a single bold line specifies an Asynchronous Vault relationship, and a bold line and non-bold line specifies a Synchronous relationship. The table below indicates if the relationship is StrictSync or Sync.

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
- If the volume is a FlexGroup volume
- If the volume is in a synchronous protection relationship
- When the volume ID is unknown, for example, when you have a intercluster relationship and the
 destination cluster has not yet been discovered
 Clicking another volume in the topology selects and displays information for that volume. A question
 mark (?) in the upper-left corner of a volume indicates that either the volume is missing or that it has
 not yet been discovered. It might also indicate that the capacity information is missing. Positioning your
 cursor over the question mark displays additional information, including suggestions for remedial
 action.

The topology displays information about volume capacity, lag, Snapshot copies, and last successful data transfer if it conforms to one of several common topology templates. If a topology does not conform to one of those templates, information about volume lag and last successful data transfer is displayed in a relationship table under the topology. In that case, the highlighted row in the table indicates the selected volume, and, in the topology view, bold lines with a blue dot indicate the relationship between the selected volume and its source volume.

Topology views include the following information:

Capacity

Displays the total amount of capacity used by the volume. Positioning your cursor over a volume in the topology displays the current warning and critical threshold settings for that volume in the Current Threshold Settings dialog box. You can also edit the threshold settings by clicking the **Edit Thresholds**

link in the Current Threshold Settings dialog box. Clearing the **Capacity** check box hides all capacity information for all volumes in the topology.

Lag

Displays the lag duration and the lag status of the incoming protection relationships. Clearing the **Lag** check box hides all lag information for all volumes in the topology. When the **Lag** check box is dimmed, then the lag information for the selected volume is displayed in the relationship table below the topology, as well as the lag information for all related volumes.

Snapshot

Displays the number of Snapshot copies available for a volume. Clearing the **Snapshot** check box hides all Snapshot copy information for all volumes in the topology. Clicking a Snapshot copy icon (

) displays the Snapshot copy list for a volume. The Snapshot copy count displayed next to the icon is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

Last Successful Transfer

Displays the amount, duration, time, and date of the last successful data transfer. When the **Last Successful Transfer** check box is dimmed, then the last successful transfer information for the selected volume is displayed in the relationship table below the topology, as well as the last successful transfer information for all related volumes.

History

Displays in a graph the history of incoming SnapMirror and SnapVault protection relationships for the selected volume. There are three history graphs available: incoming relationship lag duration, incoming relationship transfer duration, 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. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Protection history graphs display the following information:

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 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.



This chart is not available for volumes that are in synchronous protection relationships.

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.



This chart is not available for volumes that are in synchronous protection relationships.

History area

The History area displays graphs that provide information about the capacity and space reservations of the selected volume. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Graphs might be empty and the message No data found displayed when the data or the state of the volume remains unchanged for a period of time.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends—for example, if the volume usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Volume Capacity Used

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

Volume Capacity Used vs Total

Displays the trend in how volume capacity is used based on the usage history, as well as the used capacity, total capacity, and details of the space savings from deduplication and compression, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Volume Capacity Used (%)

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

Snapshot Capacity Used (%)

Displays the Snapshot reserve and Snapshot warning threshold as line graphs, and the capacity used by the Snapshot copies as an area graph, in percentage, on the vertical (y) axis. The Snapshot overflow is represented with different colors. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Snapshot Reserve legend, the Snapshot Reserve graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity

Displays the severity of the event.

Event

Displays the event name.

Triggered Time

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp when the event was generated is displayed.

Related Annotations pane

The Related Annotations pane enables you to view annotation details associated with the selected volume. The details include the annotation name and the annotation values that are applied to the volume. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view and navigate to the SVMs, aggregates, qtrees, LUNs, and Snapshot copies that are related to the volume:

Storage Virtual Machine

Displays the capacity and the health status of the SVM that contains the selected volume.

Aggregate

Displays the capacity and the health status of the aggregate that contains the selected volume. For FlexGroup volumes, the number of aggregates that comprise the FlexGroup is listed.

· Volumes in the Aggregate

Displays the number and capacity of all the volumes that belong to the parent aggregate of the selected volume. The health status of the volumes is also displayed, based on the highest severity level. For example, if an aggregate contains ten volumes, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical. This component does not appear for FlexGroup volumes.

Qtrees

Displays the number of qtrees that the selected volume contains and the capacity of qtrees with quota that the selected volume contains. The capacity of the qtrees with quota is displayed in relation to the volume data capacity. The health status of the qtrees is also displayed, based on the highest severity level. For example, if a volume has ten qtrees, five with Warning status and the remaining five with Critical status, then the status displayed is Critical.

NFS Exports

Displays the number and status of the NFS exports associated with the volume.

CIFS Shares

Displays the number and status of the CIFS shares.

• LUNs

Displays the number and total size of all the LUNs in the selected volume. The health status of the LUNs is also displayed, based on the highest severity level.

User and Group Quotas

Displays the number and status of the user and user group quotas associated with the volume and its qtrees.

FlexClone Volumes

Displays the number and capacity of all the cloned volumes of the selected volume. The number and capacity are displayed only if the selected volume contains any cloned volumes.

Parent Volume

Displays the name and capacity of the parent volume of a selected FlexClone volume. The parent volume is displayed only if the selected volume is a FlexClone volume.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected volume.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected volume. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Health/Storage Virtual Machine details page

You can use the Health/Storage Virtual Machine details page to view detailed information about the selected SVM, such as its health, capacity, configuration, data policies, logical

interfaces (LIFs), LUNs, qtrees, and user and user group quotas. You can also view information about the related objects and related alerts for the SVM.



You can monitor only data SVMs.

Command buttons

The command buttons enable you to perform the following tasks for the selected SVM:

Switch to Performance View

Enables you to navigate to the Performance/SVM details page.

Actions

Add Alert

Enables you to add an alert to the selected SVM.

Edit Thresholds

Enables you to edit the SVM thresholds.



This button is enabled only when on the Qtrees tab, or for an SVM with Infinite Volume.

Annotate

Enables you to annotate the selected SVM.

View Storage Virtual Machines

Enables you to navigate to the Health/Storage Virtual Machines inventory page.

Health tab

The Health tab displays detailed information about data availability, data capacity, and protection issues of various objects such as volumes, aggregates, NAS LIFs, SAN LIFs, LUNs, protocols, services, NFS exports, and CIFS shares.

You can click the graph of an object to view the filtered list of objects. For example, you can click the volume capacity graph that displays warnings to view the list of volumes that have capacity issues with severity as warning.

Availability Issues

Displays, as a graph, the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the SVM. For example, information is displayed about the NAS LIFs and the SAN LIFs that are down and volumes that are offline.

You can also view information about the related protocols and services that are currently running, and the number and status of NFS exports and CIFS shares.

If the selected SVM is an SVM with Infinite Volume, you can view availability details about the Infinite Volume.

· Capacity Issues

Displays, as a graph, the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the SVM. For example, information is displayed about aggregates that are likely to breach the set threshold values.

If the selected SVM is an SVM with Infinite Volume, you can view capacity details about the Infinite Volume.

Protection Issues

Provides a quick overview of SVM protection-related health by displaying, as a graph, the total number of relationships, including relationships that have protection issues and relationships that do not have any protection-related issues. When unprotected volumes exist, clicking on the link takes you to the Health/Volumes inventory page where you can view a filtered list of the unprotected volumes on the SVM. The colors in the graph represent the different severity levels of the issues. Clicking a graph takes you to the Protection/Volume Relationships page, where you can view a filtered list of protection relationship details. The information below the graph provides details about protection issues that can impact or have already impacted the protection of data in the SVM. For example, information is displayed about volumes that have a Snapshot copy reserve that is almost full or about SnapMirror relationship lag issues.

If the selected SVM is a repository SVM, the Protection area does not display.

Capacity tab

The Capacity tab displays detailed information about the data capacity of the selected SVM.

The following information is displayed for an SVM with FlexVol volume or FlexGroup volume:

Capacity

The Capacity area displays details about the used and available capacity allocated from all volumes:

Total Capacity

Displays the total capacity (in MB, GB, and so on) of the SVM.

Used

Displays the space used by data in the volumes that belong to the SVM.

Guaranteed Available

Displays the guaranteed available space for data that is available for volumes in the SVM.

Unquaranteed

Displays the available space remaining for data that is allocated for thinly provisioned volumes in the SVM.

Volumes with Capacity Issues

The Volumes with Capacity Issues list displays, in tabular format, details about the volumes that have capacity issues:

Status

Indicates that the volume has a capacity-related issue of an indicated severity.

You can move the pointer over the status to view more information about the capacity-related event or events generated for the volume.

If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use the **View Details** button to view more information about the event.

If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.



A volume can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a volume has two events with severities of Error and Warning, only the Error severity is displayed.

Volume

Displays the name of the volume.

Used Data Capacity

Displays, as a graph, information about the volume capacity usage (in percentage).

Days to Full

Displays the estimated number of days remaining before the volume reaches full capacity.

Thin Provisioned

Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.

Aggregates

For FlexVol volumes, displays the name of the aggregate that contains the volume. For FlexGroup volumes, displays the number of aggregates that are used in the FlexGroup.

The following information is displayed for an SVM with Infinite volume:

Capacity

Displays the following capacity-related details:

- Percentage of used and free data capacity
- Percentage of used and free Snapshot capacity

Snapshot Overflow

Displays the data space that is consumed by the Snapshot copies.

Used

Displays the space used by data in the SVM with Infinite Volume.

Warning

Indicates that the space in the SVM with Infinite Volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

Error

Indicates that the space in the SVM with Infinite Volume if full. If this threshold is breached, the Space Full event is generated.

Other Details

Total Capacity

Displays the total capacity in the SVM with Infinite Volume.

Data Capacity

Displays used data capacity, available data capacity, and Snapshot overflow capacity details of the SVM with Infinite Volume.

Snapshot Reserve

Displays the used and free details of the Snapshot reserve.

System Capacity

Displays the used system capacity and available system capacity in the SVM with Infinite Volume.

Thresholds

Displays the nearly full and full thresholds of the SVM with Infinite Volume.

Storage Class Capacity Details

Displays information about the capacity usage in your storage classes. This information is displayed only if you have configured storage classes for your SVM with Infinite Volume.

Storage Virtual Machine Storage Class Thresholds

Displays the following thresholds (in percentage) of your storage classes:

Nearly Full Threshold

Specifies the percentage at which a storage class in an SVM with Infinite Volume is considered to be nearly full.

Full Threshold

Specifies the percentage at which the storage class in an SVM with Infinite Volume is considered full.

Snapshot Usage Limit

Specifies the limit, in percentage, on the space reserved for Snapshot copies in the storage class.

Configuration tab

The Configuration tab displays configuration details about the selected SVM, such as its cluster, root volume, the type of volumes it contains (Infinite Volume or FlexVol volumes), and the policies created on the SVM:

Overview

Cluster

Displays the name of the cluster to which the SVM belongs.

Allowed Volume Type

Displays the type of volumes that can be created in the SVM. The type can be InfiniteVol, FlexVol, or FlexVol/FlexGroup.

Root Volume

Displays the name of the root volume of the SVM.

Allowed Protocols

Displays the type of protocols that can be configured on the SVM. Also, indicates if a protocol is up (), down (), or is not configured ().

Data LIFs

NAS

Displays the number of NAS LIFs that are associated with the SVM. Also, indicates if the LIFs are up () or down ().

∘ SAN

Displays the number of SAN LIFs that are associated with the SVM. Also, indicates if the LIFs are up () or down ().

• FC-NVMe

Displays the number of FC-NVMe 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 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, FC-NVMe, and FlexCache. For Infinite Volume, the SAN protocols are not applicable.

Role

Displays the LIF role. The roles can be Data or Management.

Home Port

Displays the physical port to which the LIF was originally associated.

Current Port

Displays the physical port to which the LIF is currently associated. If the LIF is migrated, the current port might be different from the home port.

Port Set

Displays the port set to which the LIF is mapped.

Failover Policy

Displays the failover policy that is configured for the LIF. For NFS, CIFS, and FlexCache LIFs, the default failover policy is Next Available. Failover policy is not applicable for FC and iSCSI LIFs.

Routing Groups

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

Failover Group

Displays the name of the failover group.

Qtrees tab

The Qtrees tab displays details about qtrees and their quotas. You can click the **Edit Thresholds** button if you want to edit the health threshold settings for qtree capacity for one or more qtrees.

Use the **Export** button to create a comma-separated values (.csv) file containing the details of all the monitored qtrees. When exporting to a CSV file you can choose to create a qtrees report for the current SVM, for all SVMs in the current cluster, or for all SVMs for all clusters in your data center. Some additional qtrees fields appear in the exported CSV file.



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 (\bigotimes), Error (\bigoplus), Warning (\triangle), or Normal (\bigotimes).

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.



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 gtree.

Cluster

Displays the name of the cluster containing the gtree. Appears only in the exported CSV file.

Storage Virtual Machine

Displays the storage virtual machine (SVM) name containing the qtree. Appears only in the exported CSV file.

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.

· Quota Type

Specifies if the quota is for a user, user group, or a gtree. Appears only in the exported CSV file.

· User or Group

Displays the name of the user or user group. There will be multiple rows for each user and user group. When the quota type is qtree or if the quota is not set, then the column is empty. Appears only in the exported CSV file.

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 in the grid page and the field is blank in the CSV export data.

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 gtree 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. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then "Not applicable" is displayed in the grid page and the field is blank in the CSV export data.

File Hard Limit

Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as "Unlimited" for the following conditions: if the quota is set without a file hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

File Soft Limit

Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as "Unlimited" for the following conditions: if the quota is set without a file soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

User and Group Quotas tab

Displays details about the user and user group quotas for the selected SVM. You can view information such as the status of the quota, name of the user or user group, soft and hard limits set on the disks and files, amount

of disk space and number of files used, and the disk threshold value. You can also change the email address associated with a user or user group.

Edit Email Address command button

Opens the Edit Email Address dialog box, which displays the current email address of the selected user or user group. You can modify the email address. If the Edit Email Address field is blank, the default rule is used to generate an email address for the selected user or user group.

If more than one user has the same quota, the names of the users are displayed as comma-separated values. Also, the default rule is not used to generate the email address; therefore, you must provide the required email address for notifications to be sent.

· Configure Email Rules command button

Enables you to create or modify rules to generate an email address for the user or user group quotas that are configured on the SVM. A notification is sent to the specified email address when there is a quota breach.

Status

Displays the current status of the quota. The status can be Critical (\bigotimes), Warning (\bigwedge), or Normal (\bigotimes).

You can move the pointer over the status icon to view more information about the event or events generated for the guota.

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.



A quota can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a quota has two events with severities of Error and Warning, only the Error severity is displayed.

User or Group

Displays the name of the user or user group. If more than one user has the same quota, the names of the users are displayed as comma-separated values.

The value is displayed as "Unknown" when ONTAP does not provide a valid user name because of SecD errors.

Type

Specifies if the quota is for a user or a user group.

Volume or Qtree

Displays the name of the volume or gtree on which the user or user group guota is specified.

You can move the pointer over the name of the volume or qtree to view more information about the volume or qtree.

Disk Used %

Displays the percentage of disk space used. The value is displayed as "Not applicable" if the quota is set without a disk hard limit.

Disk Hard Limit

Displays the maximum amount of disk space allocated for the quota. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as "Unlimited" if the quota is set without a disk hard limit.

Disk Soft Limit

Displays the amount of disk space allocated for the quota before a warning event is generated. The value is displayed as "Unlimited" if the quota is set without a disk soft limit. By default, this column is hidden.

Disk Threshold

Displays the threshold value set on the disk space. The value is displayed as "Unlimited" if the quota is set without a disk threshold limit. By default, this column is hidden.

Files Used %

Displays the percentage of files used in the qtree. The value is displayed as "Not applicable" if the quota is set without a file hard limit.

File Hard Limit

Displays the hard limit for the number of files permitted on the quota. The value is displayed as "Unlimited" if the quota is set without a file hard limit.

• File Soft Limit

Displays the soft limit for the number of files permitted on the quota. The value is displayed as "Unlimited" if the quota is set without a file soft limit. By default, this column is hidden.

Email Address

Displays the email address of the user or user group to which notifications are sent when there is a breach in the quotas.

NFS Exports tab

The NFS Exports tab displays information about NFS exports such as its status, the path associated with the volume (Infinite Volumes, FlexGroup volumes, or FlexVol volumes), access levels of clients to the NFS exports, and the export policy defined for the volumes that are exported. NFS exports will not be displayed in the following conditions: if the volume is not mounted or if the protocols associated with the export policy for the volume do not contain NFS exports.

Use the **Export** button to create a comma-separated values (.csv) file containing the details of all the monitored NFS exports. When exporting to a CSV file you can choose to create an NFS exports report for the current SVM, for all SVMs in the current cluster, or for all SVMs for all clusters in your data center. Some

additional export policy fields appear in the exported CSV file.

Status

Displays the current status of the NFS export. The status can be Error (1) or Normal (2).

Junction Path

Displays the path to which the volume is mounted. If an explicit NFS exports policy is applied to a qtree, the column displays the path of the volume through which the qtree can be accessed.

Junction Path Active

Displays whether the path to access the mounted volume is active or inactive.

Volume or Qtree

Displays the name of the volume or qtree to which the NFS export policy is applied. For Infinite Volumes, the name of the SVM with the Infinite Volume is displayed. If an NFS export policy is applied to a qtree in the volume, the column displays both the names of the volume and the qtree.

You can click the link to view details about the object in the respective details page. If the object is a qtree, links are displayed for both the qtree and the volume.

Cluster

Displays the name of the cluster. Appears only in the exported CSV file.

Storage Virtual Machine

Displays the name of the SVM with NFS export policies. Appears only in the exported CSV file.

Volume State

Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, or Mixed.

Offline

Read or write access to the volume is not allowed.

Online

Read and write access to the volume is allowed.

Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Mixed

The constituents of a FlexGroup volume are not all in the same state.

· Security Style

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified,

NTFS, or Mixed.

UNIX (NFS clients)

Files and directories in the volume have UNIX permissions.

Unified

Files and directories in the volume have a unified security style.

NTFS (CIFS clients)

Files and directories in the volume have Windows NTFS permissions.

Mixed

Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

UNIX Permission

Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.

Export Policy

Displays the rules that define the access permission for volumes that are exported. You can click the link to view details about the rules associated with the export policy such as the authentication protocols and the access permission.

When you generate a report for the NFS Exports page, all rules that belong to the export policy are exported to the CSV file. For example, if there are two rules in the export policy, you will see only one row in the NFS Exports grid page, but the exported data will have two rows corresponding to the two rules.

Rule Index

Displays the rules associated with the export policy such as the authentication protocols and the access permission. Appears only in the exported CSV file.

Access Protocols

Displays the protocols that are enabled for the export policy rules. Appears only in the exported CSV file.

Client Match

Displays the clients that have permission to access data on the volumes. Appears only in the exported CSV file.

Read Only Access

Displays the authentication protocol used to read data on the volumes. Appears only in the exported CSV file.

Read Write Access

Displays the authentication protocol used to read or write data on the volumes. Appears only in the exported CSV file.

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.



Shares in folders are not displayed in the CIFS Shares tab.

View User Mapping command button

Launches the User Mapping dialog box.

You can view the details of user mapping for the SVM.

Show ACL command button

Launches the Access Control dialog box for the share.

You can view user and permission details for the selected share.

Status

Displays the current status of the share. The status can be Normal () or Error ().

Share Name

Displays the name of the CIFS share.

• Path

Displays the junction path on which the share is created.

Junction Path Active

Displays whether the path to access the share is active or inactive.

Containing Object

Displays the name of the containing object to which the share belongs. The containing object can be a volume or a qtree.

By clicking the link, you can view details about the containing object in the respective Details page. If the containing object is a qtree, links are displayed for both qtree and volume.

Volume State

Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, or Mixed.

· Offline

Read or write access to the volume is not allowed.

Online

Read and write access to the volume is allowed.

Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

Mixed

The constituents of a FlexGroup volume are not all in the same state.

Security

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

UNIX (NFS clients)

Files and directories in the volume have UNIX permissions.

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.

Export Policy

Displays the name of the export policy applicable to the share. If an export policy is not specified for the SVM, the value is displayed as Not Enabled.

You can click the link to view details about the rules associated with the export policy, such as access protocols and permissions. The link is disabled if the export policy is disabled for the selected SVM.

NFS Equivalent

Specifies whether there is an NFS equivalent for the share.

SAN tab

Displays details about LUNs, initiator groups, and initiators for the selected SVM. By default, the LUNs view is displayed. You can view details about the initiator groups in the Initiator Groups tab and details about initiators in the Initiators tab.

· LUNs tab

Displays details about the LUNs that belong to the selected SVM. You can view information such as the LUN name, LUN state (online or offline), the name of the file system (volume or qtree) that contains the LUN, the type of host operating system, the total data capacity and serial number of the LUN. You can also view information whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group.

You can also view the initiator groups and initiators that are mapped to the selected LUN.

Initiator Groups tab

Displays details about initiator groups. You can view details such as the name of the initiator group, the access state, the type of host operating system that is used by all the initiators in the group, and the supported protocol. When you click the link in the access state column, you can view the current access state of the initiator group.

Normal

The initiator group is connected to multiple access paths.

Single Path

The initiator group is connected to a single access path.

No Paths

There is no access path connected to the initiator group.

You can view whether initiator groups are mapped to all the LIFs or specific LIFs through a port set. When you click the count link in the Mapped LIFs column, either all LIFs are displayed or specific LIFs for a port set are displayed. LIFs that are mapped through the target portal are not displayed. The total number of initiators and LUNs that are mapped to an initiator group is displayed.

You can also view the LUNs and initiators that are mapped to the selected initiator group.

Initiators tab

Displays the name and type of the initiator and the total number of initiator groups mapped to this initiator for the selected SVM.

You can also view the LUNs and initiator groups that are mapped to the selected initiator group.

Data Policy tab

The Data Policy tab enables you to create, modify, activate, or delete one or more rules in a data policy. You can also import the data policy into the Unified Manager database and export the data policy to your computer:



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".



The file path must always start with a junction path.

Content Placement

Displays the corresponding storage class for the rule.

Rule Filter

Enables you to filter rules associated with a specific storage class listed in the list.

Action buttons

Create

Opens the Create Rule dialog box, which enables you to create a new rule for your data policy.

• Edit

Opens the Edit Rule dialog box, which enables you to modify rule properties such as directory paths, file types, and owners.

• Delete

Deletes the selected rule.

Move Up

Moves the selected rule up in the list. However, you cannot move the default rule up in the list.

Move Down

Moves the selected rule down the list. However, you cannot move the default rule down the list.

· Activate

Activates the rules and changes made to the data policy in the SVM with Infinite Volume.

Reset

Resets all changes made to the data policy configuration.

Import

Imports a data policy configuration from a file.

Export

Exports a data policy configuration to a file.

Related Devices area

The Related Devices area enables you to view and navigate to the LUNs, CIFS shares, and the user and user group quotas that are related to the qtree:

• LUNs

Displays the total number of the LUNs associated with the selected qtree.

NFS exports

Displays the total number of NFS export policies associated with the selected gtree.

CIFS Shares

Displays the total number of CIFS shares associated with the selected qtree.

User and Group Quotas

Displays the total number of the user and user group quotas associated with the selected qtree. The health status of the user and user group quotas is also displayed, based on the highest severity level.

Related Annotations pane

The Related Annotations pane enables you to view the annotation details associated with the selected SVM. Details include the annotation name and the annotation values that are applied to the SVM. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view the cluster, aggregates, and volumes that are related to the SVM:

Cluster

Displays the health status of the cluster to which the SVM belongs.

Aggregates

Displays the number of aggregates that belong to the selected SVM. The health status of the aggregates is also displayed, based on the highest severity level. For example, if an SVM contains ten aggregates, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.

Assigned Aggregates

Displays the number of aggregates that are assigned to an SVM. The health status of the aggregates is also displayed, based on the highest severity level.

Volumes

Displays the number and capacity of the volumes that belong to the selected SVM. The health status of the volumes is also displayed, based on the highest severity level. When there are FlexGroup volumes in the SVM, the count also includes FlexGroups; it does not include FlexGroup constituents.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected SVM.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected SVM. You can also add an alert by clicking the **Add Alert** link or edit an existing alert by clicking the alert name.

Health/Cluster details page

The Health/Cluster details page provides detailed information about a selected cluster, such as health, capacity, and configuration details. You can also view information about the logical interfaces (LIFs), nodes, disks, related devices, and related alerts for the cluster.

The status next to the cluster name, for example (Good), represents the communication status; whether Unified Manager can communicate with the cluster. It does not represent the failover status or overall status of the cluster.

Command buttons

The command buttons enable you to perform the following tasks for the selected cluster:

Switch to Performance View

Enables you to navigate to the Performance/Cluster details page.



Enables you to add the selected cluster to the Favorites dashboard.

Actions

- · Add Alert: Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.
- Rediscover: Initiates a manual refresh of the cluster, which enables Unified Manager to discover recent changes to the cluster.

If Unified Manager is paired with OnCommand Workflow Automation, the rediscovery operation also reacquires cached data from WFA, if any.

After the rediscovery operation is initiated, a link to the associated job details is displayed to enable tracking of the job status.

• Annotate: Enables you to annotate the selected cluster.

View Clusters

Enables you to navigate to the Health/Clusters inventory page.

Health tab

Displays detailed information about the data availability and data capacity issues of various cluster objects such as nodes, SVMs, and aggregates. Availability issues are related to the data-serving capability of the cluster objects. Capacity issues are related to the data-storing capability of the cluster objects.

You can click the graph of an object to view a filtered list of the objects. For example, you can click the SVM capacity graph that displays warnings to view a filtered list of SVMs. This list contains SVMs that have volumes or qtrees that have capacity issues with a severity level of Warning. You can also click the SVMs availability graph that displays warnings to view the list of SVMs that have availability issues with a severity level of Warning.

Availability Issues

Graphically displays the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the cluster. For example, information is displayed about disk shelves that are down and aggregates that are offline.



The data displayed for the SFO bar graph is based on the HA state of the nodes. The data displayed for all other bar graphs is calculated based on the events generated.

Capacity Issues

Graphically displays the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the cluster. For example, information is displayed about aggregates that are likely to breach the set threshold values.

Capacity tab

Displays detailed information about the capacity of the selected cluster.

Capacity

Displays the data capacity graph about the used capacity and available capacity from all allocated aggregates:

Total Capacity

Displays the total capacity of the cluster. This does not include the capacity that is assigned for parity.

Used

Displays the capacity that is used by data. This does not include the capacity that is used for parity, right-sizing, and reservation.

Available

Displays the capacity available for data.

Spares

Displays the storable capacity available for storage in all the spare disks.

Provisioned

Displays the capacity that is provisioned for all the underlying volumes.

Cloud Tier

Displays capacity details about the cloud tier for FabricPool-enabled aggregates on the cluster. A FabricPool can be either licensed or unlicensed.

Used

Displays the space used by data in configured cloud tiers.

Data graph

For an Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage, the chart displays the total data capacity that has been licensed by this cluster and the amount being used by aggregates.

For a StorageGRID, the chart displays only the total capacity being used by aggregates.

Details

Displays detailed information about the used and available capacity.

Total Capacity

Displays the total capacity of the cluster. This does not include the capacity that is assigned for parity.

Used

Displays the capacity that is used by data. This does not include the capacity that is used for parity, right-sizing, and reservation.

Available

Displays the capacity available for data.

Provisioned

Displays the capacity that is provisioned for all the underlying volumes.

Spares

Displays the storable capacity available for storage in all the spare disks.

Cloud Tier

Displays the space used by data in configured cloud tiers. For an Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage, the total data capacity that has been licensed by this cluster is also displayed.

Capacity Breakout by Disk Type

The Capacity Breakout by Disk Type area displays detailed information about the disk capacity of the various types of disks in the cluster. By clicking the disk type, you can view more information about the disk type from the Disks tab.

Total Usable Capacity

Displays the available capacity and spare capacity of the data disks.

• HDD

Graphically displays the used capacity and available capacity of all the HDD data disks in the cluster. The dotted line represents the spare capacity of the data disks in the HDD.

Flash

SSD Data

Graphically displays the used capacity and available capacity of the SSD data disks in the cluster.

SSD Cache

Graphically displays the storable capacity of the SSD cache disks in the cluster.

SSD Spare

Graphically displays the spare capacity of the SSD, data, and cache disks in the cluster.

Unassigned Disks

Displays the number of unassigned disks in the cluster.

Aggregates with Capacity Issues list

Displays in tabular format details about the used capacity and available capacity of the aggregates that have capacity risk issues.

Status

Indicates that the aggregate has a capacity-related issue of a certain severity.

You can move the pointer over the status to view more information about the event or events generated for the aggregate.

If the status of the aggregate is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event.

If the status of the aggregate is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.



An aggregate can have multiple capacity-related events of the same severity or different severities. However, only the highest severity is displayed. For example, if an aggregate has two events with severity levels of Error and Critical, only the Critical severity is displayed.

Aggregate

Displays the name of the aggregate.

Used Data Capacity

Graphically displays information about the aggregate capacity usage (in percentage).

Days to Full

Displays the estimated number of days remaining before the aggregate reaches full capacity.

Configuration tab

Displays details about the selected cluster, such as IP address, serial number, contact, and location:

Cluster Overview

Management LIF

Displays the cluster-management LIF that Unified Manager uses to connect to the cluster. The operational status of the LIF is also displayed.

Host Name or IP Address

Displays the FQDN, short name, or the IP address of the cluster-management LIF that Unified Manager uses to connect to the cluster.

• FQDN

Displays the fully qualified domain name (FQDN) of the cluster.

OS Version

Displays the ONTAP version that the cluster is running. If the nodes in the cluster are running different versions of ONTAP, then the earliest ONTAP version is displayed.

Serial Number

Displays the serial number of the cluster.

Contact

Displays details about the administrator whom you should contact in case of issues with the cluster.

Location

Displays the location of the cluster.

Remote Cluster Overview

Provides details about the remote cluster in a MetroCluster configuration. This information is displayed only for MetroCluster configurations.

Cluster

Displays the name of the remote cluster. You can click the cluster name to navigate to the details page of the cluster.

Hostname or IP Address

Displays the FQDN, short name, or IP address of the remote cluster.

Serial Number

Displays the serial number of the remote cluster.

Location

Displays the location of the remote cluster.

MetroCluster Overview

Provides details about the local cluster in a MetroCluster configuration. This information is displayed only for MetroCluster configurations.

Type

Displays whether the MetroCluster type is two-node or four-node.

Configuration

Displays the MetroCluster configuration, which can have the following values:

- Stretch Configuration with SAS cables
- Stretch Configuration with FC-SAS bridge
- Fabric Configuration with FC switches



For a four-node MetroCluster, only Fabric Configuration with FC switches is supported.

Automated Unplanned Switch Over (AUSO)

Displays whether automated unplanned switchover is enabled for the local cluster. By default, AUSO is enabled for all clusters in a two-node MetroCluster configuration in Unified Manager. You can use the command-line interface to change the AUSO setting.

Nodes

Availability

Displays the number of nodes that are up () or down () in the cluster.

OS Versions

Displays the ONTAP versions that the nodes are running as well as the number of nodes running a particular version of ONTAP. For example, 9.0 (2), 8.3 (1) specifies that two nodes are running ONTAP 9.0, and one node is running ONTAP 8.3.

Storage Virtual Machines

Availability

Displays the number of SVMs that are up () or down () in the cluster.

• LIFs

Availability

Displays the number of non-data LIFs that are up () or down () in the cluster.

Cluster-Management LIFs

Displays the number of cluster-management LIFs.

Node-Management LIFs

Displays the number of node-management LIFs.

Cluster LIFs

Displays the number of cluster LIFs.

Intercluster LIFs

Displays the number of intercluster LIFs.

Protocols

Data Protocols

Displays the list of licensed data protocols that are enabled for the cluster. The data protocols include iSCSI, CIFS, NFS, NVMe, and FC/FCoE.

Cloud Tiers

Lists the names of the cloud tiers to which this cluster is connected. It also lists the type (Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, Alibaba Cloud Object Storage, or StorageGRID), and the states of the cloud tiers (Available or Unavailable).

MetroCluster Connectivity tab

Displays the issues and connectivity status of the cluster components in the MetroCluster configuration. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.



The MetroCluster Connectivity tab is displayed only for clusters that are in a MetroCluster configuration.

You can navigate to the details page of a remote cluster by clicking the name of the remote cluster. You can also view the details of the components by clicking the count link of a component. For example, clicking the count link of the node in the cluster displays the node tab in the details page of the cluster. Clicking the count link of the disks in the remote cluster displays the disk tab in the details page of the remote cluster.



When managing an eight-node MetroCluster configuration, clicking the count link of the Disk Shelves component displays only the local shelves of the default HA pair. Also, there is no way to display the local shelves on the other HA pair.

You can move the pointer over the components to view the details and the connectivity status of the clusters in case of any issue and to view more information about the event or events generated for the issue.

If the status of the connectivity issue between components is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. The View Details button provides more information about the event.

If status of the connectivity issue between components is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

MetroCluster Replication tab

Displays the status of the data that is being replicated. You can use the MetroCluster Replication tab to ensure data protection by synchronously mirroring the data with the already peered clusters. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.



The MetroCluster Replication tab is displayed only for clusters that are in a MetroCluster configuration.

In a MetroCluster environment, you can use this tab to verify the logical connections and peering of the local cluster with the remote cluster. You can view the objective representation of the cluster components with their logical connections. This helps to identify the issues that might occur during mirroring of metadata and data.

In the MetroCluster Replication tab, local cluster provides the detailed graphical representation of the selected cluster and MetroCluster partner refers to the remote cluster.

LIFs tab

Displays details about all the non-data LIFs that are created on the selected cluster.

• LIF

Displays the name of the LIF that is created on the selected cluster.

Operational Status

Displays the operational status of the LIF, which can be Up ($^{+}$), Down ($^{-}$), or Unknown ($^{-}$). The operational status of a LIF is determined by the status of its physical ports.

Administrative Status

Displays the administrative status of the LIF, which can be Up (), Down (), or Unknown (). You can control the administrative status of a LIF when you make changes to the configuration or during maintenance. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.

IP Address

Displays the IP address of the LIF.

Role

Displays the role of the LIF. Possible roles are Cluster-Management LIFs, Node-Management LIFs, Cluster LIFs, and Intercluster LIFs.

Home Port

Displays the physical port to which the LIF was originally associated.

Current Port

Displays the physical port to which the LIF is currently associated. After LIF migration, the current port might be different from the home port.

Failover Policy

Displays the failover policy that is configured for the LIF.

Routing Groups

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

Failover Group

Displays the name of the failover group.

Nodes tab

Displays information about nodes in the selected cluster. You can view detailed information about the HA pairs, disk shelves, and ports:

HA Details

Provides a pictorial representation of the HA state and the health status of the nodes in the HA pair. The health status of the node is indicated by the following colors:

Green

The node is in a working condition.

Yellow

The node has taken over the partner node or the node is facing some environmental issues.

Red

The node is down.

You can view information about the availability of the HA pair and take required action to prevent any risks. For example, in the case of a possible takeover operation, the following message is displayed: Storage failover possible.

You can view a list of the events related to the HA pair and its environment, such as fans, power supplies, NVRAM battery, flash cards, service processor, and connectivity of disk shelves. You can also view the time when the events were triggered.

You can view other node-related information, such as the model number and the serial number.

If there are single-node clusters, you can also view details about the nodes.

Disk Shelves

Displays information about the disk shelves in the HA pair.

You can also view events generated for the disk shelves and the environmental components, and the time when the events were triggered.

Shelf ID

Displays the ID of the shelf where the disk is located.

Component Status

Displays environmental details of the disk shelves, such as power supplies, fans, temperature sensors, current sensors, disk connectivity, and voltage sensors. The environmental details are displayed as icons in the following colors:

Green

The environmental components are in working properly.

Grey

No data is available for the environmental components.

Red

Some of the environmental components are down.

State

Displays the state of the disk shelf. The possible states are Offline, Online, No status, Initialization required, Missing, and Unknown.

Model

Displays the model number of the disk shelf.

Local Disk Shelf

Indicates whether the disk shelf is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

Unique ID

Displays the unique identifier of the disk shelf.

Firmware Version

Displays the firmware version of the disk shelf.

Ports

Displays information about the associated FC, FCoE, and Ethernet ports. You can view details about the ports and the associated LIFs by clicking the port icons.

You can also view the events generated for the ports.

You can view the following port details:

Port ID

Displays the name of the port. For example, the port names can be e0M, e0a, and e0b.

Role

Displays the role of the port. The possible roles are Cluster, Data, Intercluster, Node-Management, and Undefined

Type

Displays the physical layer protocol used for the port. The possible types are Ethernet, Fibre Channel, and FCoE.

WWPN

Displays the World Wide Port Name (WWPN) of the port.

Firmware Rev

Displays the firmware revision of the FC/FCoE port.

Status

Displays the current state of the port. The possible states are Up, Down, Link Not Connected. or Unknown (?).

You can view the port-related events from the Events list. You can also view the associated LIF details, such as LIF name, operational status, IP address or WWPN, protocols, name of the SVM associated with the LIF, current port, failover policy and failover group.

Disks tab

Displays details about the disks in the selected cluster. You can view disk-related information such as the number of used disks, spare disks, broken disks, and unassigned disks. You can also view other details such as the disk name, disk type, and the owner node of the disk.

Disk Pool Summary

Displays the number of disks, which are categorized by effective types (FCAL, SAS, SATA, MSATA, SSD, Array LUN, and VMDISK), and the state of the disks. You can also view other details, such as the number of aggregate, shared disks, spare disks, broken disks, unassigned disks, and unsupported disks. If you click the effective disk type count link, disks of the selected state and effective type are displayed. For example, if you click the count link for the disk state Broken and effective type SAS, all disks with the disk state Broken and effective type SAS are displayed.

Disk

Displays the name of the disk.

RAID Groups

Displays the name of the RAID group.

Owner Node

Displays the name of the node to which the disk belongs. If the disk is unassigned, no value is displayed in this column.

State

Displays the state of the disk: Aggregate, Shared, Spare, Broken, Unassigned, Unsupported or Unknown. By default, this column is sorted to display the states in the following order: Broken, Unassigned, Unsupported, Spare, Aggregate, and Shared.

Local Disk

Displays either Yes or No to indicate whether the disk is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

Position

Displays the position of the disk based on its container type: for example, Copy, Data, or Parity. By default, this column is hidden.

Impacted Aggregates

Displays the number of aggregates that are impacted due to the failed disk. You can move the pointer over the count link to view the impacted aggregates and then click the aggregate name to view details of the aggregate. You can also click the aggregate count to view the list of impacted aggregates in the Health/Aggregates inventory page.

No value is displayed in this column for the following cases:

- For broken disks when a cluster containing such disks is added to Unified Manager
- When there are no failed disks

Storage Pool

Displays the name of the storage pool to which the SSD belongs. You can move the pointer over the storage pool name to view details of the storage pool.

Storable Capacity

Displays the disk capacity that is available for use.

Raw Capacity

Displays the capacity of the raw, unformatted disk before right-sizing and RAID configuration. By default, this column is hidden.

Type

Displays the types of disks: for example, ATA, SATA, FCAL, or VMDISK.

Effective Type

Displays the disk type assigned by ONTAP.

Certain ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and spare management. ONTAP assigns an effective disk type for each disk type.

Spare Blocks Consumed %

Displays in percentage the spare blocks that are consumed in the SSD disk. This column is blank for disks

other than SSD disks.

Rated Life Used %

Displays in percentage an estimate of the SSD life used, based on the actual SSD usage and the manufacturer's prediction of SSD life. A value greater than 99 indicates that the estimated endurance has been consumed, but may not indicate SSD failure. If the value is unknown, then the disk is omitted.

Firmware

Displays the firmware version of the disk.

RPM

Displays the revolutions per minute (RPM) of the disk. By default, this column is hidden.

Model

Displays the model number of the disk. By default, this column is hidden.

Vendor

Displays the name of the disk vendor. By default, this column is hidden.

Shelf ID

Displays the ID of the shelf where the disk is located.

Bay

Displays the ID of the bay where the disk is located.

Related Annotations pane

Enables you to view the annotation details associated with the selected cluster. The details include the annotation name and the annotation values that are applied to the cluster. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

Enables you to view device details that are associated with the selected cluster.

The details include properties of the device that is connected to the cluster such as the device type, size, count, and health status. You can click on the count link for further analysis on that particular device.

You can use MetroCluster Partner pane to obtain count and also details on the remote MetroCluster partner along with its associated cluster components such as nodes, aggregates, and SVMs. The MetroCluster Partner pane is displayed only for clusters in a MetroCluster configuration.

The Related Devices pane enables you to view and navigate to the nodes, SVMs, and aggregates that are related to the cluster:

MetroCluster Partner

Displays the health status of the MetroCluster partner. Using the count link, you can navigate further and obtain information about the health and capacity of the cluster components.

Nodes

Displays the number, capacity, and health status of the nodes that belong to the selected cluster. Capacity indicates the total usable capacity over available capacity.

Storage Virtual Machines

Displays the number of SVMs that belong to the selected cluster.

Aggregates

Displays the number, capacity, and the health status of the aggregates that belong to the selected cluster.

Related Groups pane

Enables you to view the list of groups that includes the selected cluster.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts for the selected cluster. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Health/Aggregate details page

You can use the Health/Aggregate details page to view detailed information about the selected aggregate, such as the capacity, disk information, configuration details, and events generated. You can also view information about the related objects and related alerts for that aggregate.

Command buttons



When monitoring a FabricPool-enabled aggregate, the committed and overcommitted values on this page are relevant only to the local, or performance tier, capacity. The amount of space available in the cloud tier is not reflected in the overcommitted values. Similarly, the aggregate threshold values are relevant only to the local performance tier.

The command buttons enable you to perform the following tasks for the selected aggregate:

Switch to Performance View

Enables you to navigate to the Performance/Aggregate details page.



Enables you to add the selected aggregate to the Favorites dashboard.

Actions

Add Alert

Enables you to add an alert to the selected aggregate.

· Edit Thresholds

Enables you to modify the threshold settings for the selected aggregate.

View Aggregates

Enables you to navigate to the Health/Aggregates inventory page.

Capacity tab

The Capacity tab displays detailed information about the selected aggregate, such as its capacity, thresholds, and daily growth rate.

By default, capacity events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are modified by a technical support representative, the threshold values are applied to the node root aggregate.

Capacity

Displays the data capacity graph and the Snapshot copies graph, which display capacity details about the aggregate:

Used

Displays the space used by data in the aggregate.

Overcommitted

Indicates that the space in the aggregate is overcommitted.

Warning

Indicates that the space in the aggregate is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

Error

Indicates that the space in the aggregate is full. If this threshold is breached, the Space Full event is generated.

Data graph

Displays the total data capacity and the used data capacity of the aggregate. If the aggregate is overcommitted, a flag is displayed with the overcommitted capacity.

Snapshot Copies graph

This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both of the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

Cloud Tier

Displays capacity details about the cloud tier for FabricPool-enabled aggregates. A FabricPool can be either licensed or unlicensed.

Used

Displays the space used by data in the cloud tier.

Unavailable

Displays the space in the cloud tier for an Amazon S3, Microsoft Azure Cloud FabricPool, or IBM Cloud Object Storage object that cannot be used. This space may be shared with another FabricPool-enabled aggregate.

· Data graph

For an Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage, the chart displays the total data capacity that has been licensed by this cluster, the amount being used by this aggregate, and the unusable amount from other aggregates that are using the cloud tier.

For a StorageGRID, the chart displays only the total capacity being used by this aggregate.

Details

Displays detailed information about capacity.

Total Capacity

Displays the total capacity in the aggregate.

Data Capacity

Displays the amount of space used by the aggregate (used capacity) and the amount of available space in the aggregate (free capacity).

Snapshot Reserve

Displays the used and free Snapshot capacity of the aggregate.

Overcommitted Capacity

Displays the aggregate overcommitment. Aggregate overcommitment enables you to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. When thin provisioning is in use, the total size of volumes in the aggregate can exceed the total capacity of the aggregate.



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.

· Cloud Tier

For an Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage, displays the total licensed capacity, the amount used by this aggregate, the amount used by other aggregates, and the free capacity for the cloud tier. For a StorageGRID, displays only the total capacity being used by this aggregate.

Total Cache Space

Displays the total space of the solid-state drives (SSDs) or allocation units that are added to a Flash Pool aggregate. If you have enabled Flash Pool for an aggregate but have not added any SSDs, then the cache space is displayed as 0 KB.



This field is hidden if Flash Pool is disabled for an aggregate.

Aggregate Thresholds

Displays the following aggregate capacity thresholds:

Nearly Full Threshold

Specifies the percentage at which an aggregate is nearly full.

Full Threshold

Specifies the percentage at which an aggregate is full.

Nearly Overcommitted Threshold

Specifies the percentage at which an aggregate is nearly overcommitted.

Overcommitted Threshold

Specifies the percentage at which an aggregate is overcommitted.

Other Details: Daily Growth Rate

Displays the disk space used in the aggregate if the rate of change between the last two samples continues for 24 hours.

For example, if an aggregate uses 10 GB of disk space at 2 pm and 12 GB at 6 pm, the daily growth rate (GB) for this aggregate is 2 GB.

Volume Move

Displays the number of volume move operations that are currently in progress:

Volumes Out

Displays the number and capacity of the volumes that are being moved out of the aggregate.

You can click the link to view more details, such as the volume name, aggregate to which the volume is moved, status of the volume move operation, and the estimated end time.

Volumes In

Displays the number and remaining capacity of the volumes that are being moved into the aggregate.

You can click the link to view more details, such as the volume name, aggregate from which the volume is moved, status of the volume move operation, and the estimated end time.

Estimated used capacity after volume move

Displays the estimated amount of used space (as a percentage, and in KB, MB, GB, and so on) in

the aggregate after the volume move operations are complete.

Capacity Overview - Volumes

Displays graphs that provide information about the capacity of the volumes contained in the aggregate. The amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume is displayed. When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

You can select the graph you want to view from the drop-down lists. You can sort the data displayed in the graph to display details such as the used size, provisioned size, available capacity, fastest daily growth rate, and slowest growth rate. You can filter the data based on the storage virtual machines (SVMs) that contain the volumes in the aggregate. You can also view details for thinly provisioned volumes. You can view the details of specific points on the graph by positioning your cursor over the area of interest. By default, the graph displays the top 30 filtered volumes in the aggregate.

Disk Information tab

Displays detailed information about the disks in the selected aggregate, including the RAID type and size, and the type of disks used in the aggregate. The tab also graphically displays the RAID groups, and the types of disks used (such as SAS, ATA, FCAL, SSD, or VMDISK). You can view more information, such as the disk's bay, shelf, and rotational speed, by positioning your cursor over the parity disks and data disks.

• Data

Graphically displays details about dedicated data disks, shared data disks, or both. When the data disks contain shared disks, graphical details of the shared disks are displayed. When the data disks contain dedicated disks and shared disks, graphical details of both the dedicated data disks and the shared data disks are displayed.

RAID Details

RAID details are displayed only for dedicated disks.

Type

Displays the RAID type (RAID0, RAID4, RAID-DP, or RAID-TEC).

Group Size

Displays the maximum number of disks allowed in the RAID group.

Groups

Displays the number of RAID groups in the aggregate.

Disks Used

Effective Type

Displays the types of data disks (for example, ATA, SATA, FCAL, SSD, or VMDISK) in the aggregate.

Data Disks

Displays the number and capacity of the data disks that are assigned to an aggregate. Data disk details are not displayed when the aggregate contains only shared disks.

Parity Disks

Displays the number and capacity of the parity disks that are assigned to an aggregate. Parity disk details are not displayed when the aggregate contains only shared disks.

Shared Disks

Displays the number and capacity of the shared data disks that are assigned to an aggregate. Shared disk details are displayed only when the aggregate contains shared disks.

Spare Disks

Displays the disk effective type, number, and capacity of the spare data disks that are available for the node in the selected aggregate.



When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

SSD Cache

Provides details about dedicated cache SSD disks and shared cache SSD disks.

The following details for the dedicated cache SSD disks are displayed:

RAID Details

Type

Displays the RAID type (RAID0, RAID4, RAID-DP or RAID-TEC).

Group Size

Displays the maximum number of disks allowed in the RAID group.

Groups

Displays the number of RAID groups in the aggregate.

Disks Used

Effective Type

Indicates that the disks used for cache in the aggregate are of type SSD.

Data Disks

Displays the number and capacity of the data disks that are assigned to an aggregate for cache.

Parity Disks

Displays the number and capacity of the parity disks that are assigned to an aggregate for cache.

Spare Disks

Displays the disk effective type, number, and capacity of the spare disks that are available for the node in the selected aggregate for cache.



When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

Provides the following details for the shared cache:

Storage Pool

Displays the name of the storage pool. You can move the pointer over the storage pool name to view the following details:

Status

Displays the status of the storage pool, which can be healthy or unhealthy.

Total Allocations

Displays the total allocation units and the size in the storage pool.

Allocation Unit Size

Displays the minimum amount of space in the storage pool that can be allocated to an aggregate.

Disks

Displays the number of disks used to create the storage pool. If the disk count in the storage pool column and the number of disks displayed in the Disk Information tab for that storage pool do not match, then it indicates that one or more disks are broken and the storage pool is unhealthy.

Used Allocation

Displays the number and size of the allocation units used by the aggregates. You can click the aggregate name to view the aggregate details.

Available Allocation

Displays the number and size of the allocation units available for the nodes. You can click the node name to view the aggregate details.

Allocated Cache

Displays the size of the allocation units used by the aggregate.

Allocation Units

Displays the number of allocation units used by the aggregate.

Disks

Displays the number of disks contained in the storage pool.

Details

Storage Pool

Displays the number of storage pools.

Total Size

Displays the total size of the storage pools.

Cloud Tier

Displays the name of the cloud tier, if you have configured a FabricPool-enabled aggregate, and shows the total licensed capacity for Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, or Alibaba Cloud Object Storage objects.

Configuration tab

The Configuration tab displays details about the selected aggregate, such as its cluster node, block type, RAID type, RAID size, and RAID group count:

Overview

Node

Displays the name of the node that contains the selected aggregate.

· Block Type

Displays the block format of the aggregate: either 32-bit or 64-bit.

RAID Type

Displays the RAID type (RAID0, RAID4, RAID-DP, RAID-TEC or Mixed RAID).

• RAID Size

Displays the size of the RAID group.

RAID Groups

Displays the number of RAID groups in the aggregate.

SnapLock Type

Displays the SnapLock Type of the aggregate.

Cloud Tier

If this is a FabricPool-enabled aggregate, the details for the object store are displayed. Some fields are different depending on the storage provider:

Name

Displays the name of the object store when it was created by ONTAP.

Object Storage Provider

Displays the name of the storage provider, for example, StorageGRID, Amazon S3, IBM Cloud Object Storage, Microsoft Azure Cloud, or Alibaba Cloud Object Storage.

Object Store Name (FQDN) or Server name

Displays the FQDN of the object store.

· Access Key or Account

Displays the access key or account for the object store.

Bucket Name or Container Name

Displays the bucket or container name of the object store.

• SSL

Displays whether SSL encryption is enabled for the object store.

History area

The History area displays graphs that provide information about the capacity of the selected aggregate. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if the aggregate usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

Aggregate Capacity Used (%)

Displays the used capacity in the aggregate and the trend in how aggregate capacity is used based on the usage history as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Capacity Used legend, the Capacity Used graph line is hidden.

Aggregate Capacity Used vs Total Capacity

Displays the trend in how aggregate capacity is used based on the usage history, as well as the used capacity and the total capacity, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

Aggregate Capacity Used (%) vs Committed (%)

Displays the trend in how aggregate capacity is used based on the usage history, as well as the committed space as line graphs, as a percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Space Committed legend, the Space Committed graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

Severity

Displays the severity of the event.

Event

Displays the event name.

Triggered Time

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp for when the event was generated is displayed.

Related Devices pane

The Related Devices pane enables you to view the cluster node, volumes, and disks that are related to the aggregate:

Node

Displays the capacity and the health status of the node that contains the aggregate. Capacity indicates the total usable capacity over available capacity.

Aggregates in the Node

Displays the number and capacity of all the aggregates in the cluster node that contains the selected aggregate. The health status of the aggregates is also displayed, based on the highest severity level. For example, if a cluster node contains ten aggregates, five of which display the Warning status and the remaining five of which display the Critical status, then the status displayed is Critical.

Volumes

Displays the number and capacity of FlexVol volumes and FlexGroup volumes in the aggregate; the number does not include FlexGroup constituents. The health status of the volumes is also displayed, based on the highest severity level.

Resource Pool

Displays the resource pools related to the aggregate.

Disks

Displays the number of disks in the selected aggregate.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected aggregate. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Protection/Job details page

The Protection/Job details page enables you to view status and other information about specific protection job tasks that are running, that are queued, or that have completed. You can use this information to monitor protection job progress and to troubleshoot job failures.

Job summary

The job summary displays the following information:

- Job ID
- Type
- State
- · Submitted Time
- Completed Time
- Duration

Command buttons

The command buttons enable you to perform the following tasks:

Refresh

Refreshes the task list and the properties associated with each task.

View Jobs

Returns you to the Protection/Jobs page.

Job tasks list

The Job tasks list displays in a table all the tasks associated with a specific job and the properties related to each task.

Started Time

Displays the day and time the task started. By default, the most recent tasks are displayed at the top of the column and older tasks are displayed at the bottom.

Type

Displays the type of task.

State

The state of a particular task:

Completed

The task has finished.

Queued

The task is about to run.

Running

The task is running.

Waiting

A job has been submitted and some associated tasks are waiting to be queued and executed.

Status

Displays the task status:

• Error ([[]])

The task failed.

∘ Normal (

The task succeeded.

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.

Adding users

You can add local users or database users by using the Management/Users page. You can also add remote users or groups that belong to an authentication server. You can assign roles to these users and, based on the privileges of the roles, users can manage the storage objects and data with Unified Manager, or view the data in a database.

Before you begin

- You must have the OnCommand Administrator role.
- To add a remote user or group, you must have enabled remote authentication and configured your authentication server.
- If you plan to configure SAML authentication so that an identity provider (IdP) authenticates users accessing the graphical interface, make sure these users are defined as "remote" users.

Access to the UI is not allowed for users of type "local" or "maintenance" when SAML authentication is enabled.

About this task

If you add a group from Windows Active Directory, then all direct members and nested subgroups can authenticate to Unified Manager, unless nested subgroups are disabled. If you add a group from OpenLDAP or other authentication services, then only the direct members of that group can authenticate to Unified Manager.

Steps

- 1. In the toolbar, click , and then click **Users** in the left Management menu.
- 2. On the Management/Users page, click Add.
- 3. In the **Add User** dialog box, select the type of user that you want to add, and enter the required information.

When entering the required user information, you must specify an email address that is unique to that user. You must avoid specifying email addresses that are shared by multiple users.

4. Click Add.

Definitions of user roles

The maintenance user or OnCommand administrator assigns a role to every user. Each role contains certain privileges. The scope of activities that you can perform in Unified Manager depends on the role you are assigned and which privileges the role contains.

Unified Manager includes the following predefined user roles:

Operator

Views storage system information and other data collected by Unified Manager, including histories and capacity trends. This 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. This role enables the storage administrator to configure thresholds and to create alerts and other storage management-specific options and policies.

OnCommand Administrator

Configures settings unrelated to storage management. This role enables the management of users, security certificates, database access, and administrative options, including authentication, SMTP, networking, and AutoSupport.



When Unified Manager is installed on Linux systems, the initial user with the OnCommand Administrator role is automatically named "umadmin".

Integration Schema

This role enables read-only access to Unified Manager database views for integrating Unified Manager with OnCommand Workflow Automation (WFA).

· Report Schema

This role enables read-only access to reporting and other database views directly from the Unified Manager database. The databases that can be viewed include:

- netapp model view
- netapp performance
- ocum
- ocum report
- ocum report birt
- opm
- scalemonitor

Definitions of user types

A user type specifies the kind of account the user holds and includes remote users, remote groups, local users, database users, and maintenance users. Each of these types has its own role, which is assigned by a user with the role of OnCommand Administrator.

Unified Manager user types are as follows:

Maintenance user

Created during the initial configuration of Unified Manager. The maintenance user then creates additional users and assigns roles. The maintenance user is also the only user with access to the maintenance console. When Unified Manager is installed on a Red Hat Enterprise Linux or CentOS system, the maintenance user is given the user name "umadmin."

Local user

Accesses the Unified Manager UI and performs functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.

· Remote group

A group of users that access the Unified Manager UI using the credentials stored on the authentication server. The name of this account should match the name of a group stored on the authentication server. All users within the remote group are given access to the Unified Manager UI using their individual user credentials. Remote groups can perform functions according to their assigned roles.

· Remote user

Accesses the Unified Manager UI using the credentials stored on the authentication server. A remote user performs functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.

· Database user

Has read-only access to data in the Unified Manager database, has no access to the Unified Manager web interface or the maintenance console, and cannot execute API calls.

Unified Manager user roles and capabilities

Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

The following table displays the functions that each user role can perform:

Function	Operator	Storage Administrator	OnCommand Administrator	Integration Schema	Report Schema
View storage system information	•	•	•	•	•
View other data, such as histories and capacity trends	•	•	•	•	•
View, assign, and resolve events	•	•	•		
View storage service objects, such as SVM associations and resource pools	•	•	•		
View threshold policies	•	•	•		

Function	Operator	Storage Administrator	OnCommand Administrator	Integration Schema	Report Schema
Manage storage service objects, such as SVM associations and resource pools		•	•		
Define alerts		•	•		
Manage storage management options		•	•		
Manage storage management policies		•	•		
Manage users			•		
Manage administrative options			•		
Define threshold policies			•		
Manage database access			•		
Manage integration with WFA and provide access to the database views				•	
Provide read- only access to reporting and other database views					•
Schedule and save reports	•	•	•		

Function	Operator	Storage Administrator	OnCommand Administrator	Integration Schema	Report Schema
Import and delete imported reports			•		

Supported Unified Manager CLI commands

As a storage administrator you can use the CLI commands to perform queries on the storage objects; for example, on clusters, aggregates, volumes, qtrees, and LUNs. You can use the CLI commands to query the Unified Manager internal database and the ONTAP database. You can also use CLI commands in scripts that are executed at the beginning or end of an operation or are executed when an alert is triggered.

All commands must be preceded with the command um cli login and a valid user name and password for authentication.

CLI command	Description	Output
<pre>um run cmd [-t <timeout>] <cluster> <command/></cluster></timeout></pre>	The simplest way to run a command on one or more hosts. Mainly used for alert scripting to get or perform an operation on ONTAP. The optional timeout argument sets a maximum time limit (in seconds) for the command to complete on the client. The default is 0 (wait forever).	As received from ONTAP.
um run query <sql command=""></sql>	Executes an SQL query. Only queries that read from the database are allowed. Any update, insert, or delete operations are not supported.	Results are displayed in a tabular form. If an empty set is returned, or if there is any syntax error or bad request, it displays the appropriate error message.

CLI command	Description	Output
<pre>um datasource add -u</pre>	Adds a datasource to the list of managed storage systems. A datasource describes how connections to storage systems are made. The options -u (username) and -P (password) must be specified when adding a datasource. The option -t (protocol) specifies the protocol used to communicate with the cluster (http or https). If the protocol is not specified, then both protocols will be attempted The option -p (port) specifies the port used to communicate with the cluster. If the port is not specified, then the default value of the appropriate protocol will be attempted. This command can be executed only by the storage admin.	Prompts for the user accept the certificate and prints the corresponding message.
um datasource list [<datasource-id>]</datasource-id>	Displays the datasources for managed storage systems.	Displays the following values in tabular format: ID Address Port, Protocol Acquisition Status, Analysis Status, Communication status, Acquisition Message, and Analysis Message.
<pre>um datasource modify [-h <hostname-or-ip>] [-u <username>] [-P <password>] [-t <protocol>] [-p <port>] <datasource-id></datasource-id></port></protocol></password></username></hostname-or-ip></pre>	Modifies one or more datasource options. Can be executed only by the storage admin.	Displays the corresponding message.
um datasource remove <datasource-id></datasource-id>	Removes the datasource from Unified Manager.	Displays the corresponding message.
<pre>um option list [<option>]</option></pre>	Lists options.	Displays the following values in tabular format: Name, Value, Default Value, and Requires Restart.
<pre>um option set <option- name="">=<option-value> [<option-name>=<option- value="">]</option-></option-name></option-value></option-></pre>	Sets one or more options. The command can be executed only by the storage admin.	Displays the corresponding message.

CLI command	Description	Output
um version	Displays the Unified Manager software version .	Version ("7.0")
<pre>um lun list [-q] [-ObjectType <object-id>]</object-id></pre>	Lists the LUNs after filtering on the specified objectq is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: um lun list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the LUNs within the cluster with ID 1.	Displays the following values in tabular format: ID and LUN path.
<pre>um svm list [-q] [-ObjectType <object-id>]</object-id></pre>	Lists the SVMs after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: um svm list-cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the SVMs within the cluster with ID 1.	Displays the following values in tabular format: Name and Cluster ID.
<pre>um qtree list [-q] [-ObjectType <object-id>]</object-id></pre>	Lists the qtrees after filtering on the specified objectq is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, svm. For example: um qtree list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the qtrees within the cluster with ID 1.	Displays the following values in tabular format: Qtree ID and Qtree Name.

CLI command	Description	Output
<pre>um disk list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the disks after filtering on the specified object. ObjectType can be disk, aggr, node, cluster. For example: um disk list-cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the disks within the cluster with ID 1.	Displays the following values in tabular format ObjectType and object-id.
<pre>um cluster list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the clusters after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, lun, qtree, volume, quota, svm. For example:um cluster list -aggr 1 In this example, "-aggr" is the objectType and "1" is the objectId. The command lists the cluster to which the aggregate with ID 1 belongs.	Displays the following values in tabular format: Name, Full Name, Serial Number, Datasource Id, Last Refresh Time, and Resource Key.
<pre>um cluster node list [-q] [-ObjectType <object-id>]</object-id></pre>	Lists the cluster nodes after filtering on the specified object. ObjectType can be disk, aggr, node, cluster. For example: um cluster node list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the nodes within the cluster with ID 1.	Displays the following values in tabular format Name and Cluster ID.
<pre>um volume list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the volumes after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, svm, aggregate. For example: um volume list-cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the volumes within the cluster with ID 1.	Displays the following values in tabular format Volume ID and Volume Name.

CLI command	Description	Output
<pre>um quota user list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the quota users after filtering on the specified object. ObjectType can be qtree, cluster, volume, quota, svm. For example: um quota user list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the quota users within the cluster with ID 1.	Displays the following values in tabular format ID, Name, SID and Email.
<pre>um aggr list [-q] [- ObjectType <object-id>]</object-id></pre>	Lists the aggregates after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, volume. For example: um aggr list -cluster 1 In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the aggregates within the cluster with ID 1.	Displays the following values in tabular format Aggr ID, and Aggr Name.
um event ack <event-ids></event-ids>	Acknowledges one or more events.	Displays the corresponding message.
um event resolve <event-ids></event-ids>	Resolves one or more events.	Displays the corresponding message.
um event assign -u <username> <event-id></event-id></username>	Assigns an event to a user.	Displays the corresponding message.
<pre>um event list [-s <source/>] [-S <event- state-filter-list="">] [<event-id>]</event-id></event-></pre>	Lists the events generated by the system or user. Filters events based on source, state, and IDs.	Displays the following values in tabular format Source, Source type, Name, Severity, State, User and Timestamp.
um cli login -u <username> [-p <password></password></username>	Logs in to the CLI. The session expires after three hours from the time of login, after which the user must login again.	Displays the corresponding message.
um cli logout	Logs out of the CLI.	Displays the corresponding message.

CLI command	Description	Output
<pre>um backup restore -f <backup_file_path_and_name></backup_file_path_and_name></pre>	Restores a database backup using .7z files.	Displays the corresponding message.
um help	Displays all first level subcommands.	Displays all first level subcommands.

Using the maintenance console

You can use the maintenance console to configure network settings, to configure and manage the system on which Unified Manager is installed, and to perform other maintenance tasks that help you prevent and troubleshoot possible issues.

What functionality the maintenance console provides

The Unified Manager maintenance console enables you to maintain the settings on your Unified Manager system and to make any necessary changes to prevent issues from occurring.

Depending on the operating system on which you have installed Unified Manager, the maintenance console provides the following functions:

- Troubleshoot any issues with your virtual appliance, especially if the Unified Manager web interface is not available
- Upgrade to newer versions of Unified Manager
- Generate support bundles to send to technical support
- · Configure network settings
- · Change the maintenance user password
- · Connect to an external data provider to send performance statistics
- Change the performance data collection internal
- Restore the Unified Manager database and configuration settings from a previously backed up version.

What the maintenance user does

The maintenance user is created during the installation of Unified Manager on a Red Hat Enterprise Linux or CentOS system. The maintenance user name is the "umadmin" user. The maintenance user has the OnCommand administrator role in the web UI, and that user can create subsequent users and assign them roles.

The maintenance user, or umadmin user, can also access the Unified Manager maintenance console.

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.

Maintenance console menus

The maintenance console consists of different menus that enable you to maintain and manage special features and configuration settings of the Unified Manager server.

Depending on the operating system on which you have installed Unified Manager, the maintenance console consists of the following menus:

- Upgrade Unified Manager (VMware only)
- Network Configuration (VMware only)
- System Configuration (VMware only)
- Support/ Diagnostics
- · Reset Server Certificate
- · External Data Provider
- Performance Polling Interval Configuration

Network Configuration menu

The Network Configuration menu enables you to manage the network settings. You should use this menu when the Unified Manager user interface is not available.



This menu is not available if Unified Manager is installed on Red Hat Enterprise Linux, CentOS, or on Microsoft Windows.

The following menu choices are available.

Display IP Address Settings

Displays the current network settings for the virtual appliance, including the IP address, network, broadcast address, netwask, gateway, and DNS servers.

Change IP Address Settings

Enables you to change any of the network settings for the virtual appliance, including the IP address, netmask, gateway, or DNS servers. If you switch your network settings from DHCP to static networking using the maintenance console, you cannot edit the host name. You must select **Commit Changes** for the changes to take place.

Display Domain Name Search Settings

Displays the domain name search list used for resolving host names.

Change Domain Name Search Settings

Enables you to change the domain names for which you want to search when resolving host names. You must select **Commit Changes** for the changes to take place.

Display Static Routes

Displays the current static network routes.

Change Static Routes

Enables you to add or delete static network routes. You must select **Commit Changes** for the changes to take place.

Add Route

Enables you to add a static route.

Delete Route

Enables you to delete a static route.

Back

Takes you back to the Main Menu.

• Exit

Exits the maintenance console.

Disable Network Interface

Disables any available network interfaces. If only one network interface is available, you cannot disable it. You must select **Commit Changes** for the changes to take place.

Enable Network Interface

Enables available network interfaces. You must select **Commit Changes** for the changes to take place.

Commit Changes

Applies any changes made to the network settings for the virtual appliance. You must select this option to enact any changes made, or the changes do not occur.

· Ping a Host

Pings a target host to confirm IP address changes or DNS configurations.

Restore to Default Settings

Resets all settings to the factory default. You must select Commit Changes for the changes to take place.

Back

Takes you back to the Main Menu.

Exit

Exits the maintenance console.

System Configuration menu

The System Configuration menu enables you to manage your virtual appliance by providing various options, such as viewing the server status, and rebooting and shutting down the virtual machine.



The System Configuration menu is not available if Unified Manager is installed on Red Hat Enterprise Linux, CentOS, or Microsoft Windows.

The following menu choices are available:

Display Server Status

Displays the current server status. Status options include Running and Not Running.

If the server is not running, you might need to contact technical support.

Reboot Virtual Machine

Reboots the virtual machine, stopping all services. After rebooting, the virtual machine and services restart.

Shut Down Virtual Machine

Shuts down the virtual machine, stopping all services.

You can select this option only from the virtual machine console.

Change < logged in user > User Password

Changes the password of the user that is currently logged in, which can only be the maintenance user.

· Increase Data Disk Size

Increases the size of the data disk (disk 3) in the virtual machine.

· Increase Swap Disk Size

Increases the size of the swap disk (disk 2) in the virtual machine.

Change Time Zone

Changes the time zone to your location.

Change NTP Server

Changes the NTP Server settings, such as IP address or fully qualified domain name (FQDN).

Restore from an OCUM Backup

Restores the Unified Manager database and configuration settings from a previously backed up version.

Reset Server Certificate

Resets the server security certificate.

· Change hostname

Changes the name of the host on which the virtual appliance is installed.

Back

Exits the System Configuration menu and returns to the Main Menu.

Exit

Exits the maintenance console menu.

Support and Diagnostics menu

The Support and Diagnostics menu enables you to generate a support bundle.

The following menu option is available:

Generate Support Bundle

Enables you to create a 7-Zip file containing full diagnostic information in the diagnostic user's home directory. The file includes information generated by an AutoSupport message, the contents of the Unified Manager database, detailed data about the Unified Manager server internals, and verbose-level logs not normally included in AutoSupport messages.

Additional menu options

The following menu options enable you to perform various administrative tasks on the Unified Manager server.

The following menu choices are available:

· Reset Server Certificate

Regenerates the HTTPS server certificate.

You can regenerate the server certificate in the Unified Manager GUI by clicking > HTTPS Certificate > Regenerate HTTPS Certificate.

· Disable SAML authentication

Disables SAML authentication so that the identity provider (IdP) no longer provides sign-on authentication for users accessing the Unified Manager GUI. This console option is typically used when an issue with the IdP server or SAML configuration blocks users from accessing the Unified Manager GUI.

External Data Provider

Provides options for connecting Unified Manager to an external data provider. After you establish the connection, performance data is sent to an external server so that storage performance experts can chart the performance metrics using third-party software. The following options are displayed:

 Display Server Configuration--Displays the current connection and configuration settings for an external data provider.

- Add / Modify Server Connection--Enables you to enter new connection settings for an external data provider, or change existing settings.
- Modify Server Configuration--Enables you to enter new configuration settings for an external data provider, or change existing settings.
- **Delete Server Connection--**Deletes the connection to an external data provider.

After the connection is deleted, Unified Manager loses its connection to the external server.

Performance Polling Interval Configuration

Provides an option for configuring how frequently Unified Manager collects performance statistical data from clusters. The default collection interval is five minutes.

You can change this interval to ten or fifteen minutes if you find that collections from large clusters are not completing on time.

Exit

Exits the maintenance console menu.

Changing the maintenance user password on Windows

You can change the Unified Manager maintenance user password when required.

Steps

1. From the Unified Manager web UI login page, click Forgot Password.

A page is displayed that prompts for the name of the user whose password you want to reset.

2. Enter the user name and click Submit.

An email with a link to reset the password is sent to the email address that is defined for that user name.

- 3. Click the **reset password link** in the email and define the new password.
- 4. Return to the web UI and log in to Unified Manager using the new password.

After you finish

If Unified Manager is installed in a Microsoft Cluster Server (MSCS) environment, then you must change the maintenance user password on the second node of the MSCS setup. The maintenance user password for both nodes must be same.

Changing the umadmin password on Linux systems

For security reasons, you must change the default password for the Unified Manager umadmin user immediately after completing the installation process. If necessary, you can change the password again anytime later.

Before you begin

- Unified Manager must be installed on a Red Hat Enterprise Linux or CentOS Linux system.
- You must have the root user credentials for the Linux system on which Unified Manager is installed.

Steps

- 1. Log in as the root user to the Linux system on which Unified Manager is running.
- 2. Change the umadmin password: passwd umadmin

The system prompts you to enter a new password for the umadmin user.

After you finish

If Unified Manager is installed in a Veritas Cluster Server (VCS) environment, you must change the umadmin password on the second node of the VCS setup. The umadmin password for both nodes must be the same.

Adding network interfaces

You can add new network interfaces if you need to separate network traffic.

Before you begin

You must have added the network interface to the virtual appliance using vSphere.

The virtual appliance must be powered on.

About this task



You cannot perform this operation if Unified Manager is installed on Red Hat Enterprise Linux or on Microsoft Windows.

Steps

1. In the vSphere console Main Menu, select System Configuration > Reboot Operating System.

After rebooting, the maintenance console can detect the newly added network interface.

- 2. Access the maintenance console.
- 3. Select Network Configuration > Enable Network Interface.
- Select the new network interface and press Enter.

Select eth1 and press Enter.

- 5. Type **y** to enable the network interface.
- 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.

Adding disk space to the Unified Manager database directory

The Unified Manager database directory contains all of the health and performance data collected from ONTAP systems. Some circumstances may require that you increase the size of the database directory.

For example, the database directory may get full if Unified Manager is collecting data from a large number of clusters where each cluster has many nodes. You will receive a warning event when the database directory is 90% full, and a critical event when the directory is 95% full.



No additional data is collected from clusters after the directory reaches 95% full.

The steps required to add capacity to the data directory are different depending on whether Unified Manager is running on a VMware ESXi server, on a Red Hat or CentOS Linux server, or on a Microsoft Windows server.

Adding space to the data directory of the Linux host

If you allotted insufficient disk space to the /opt/netapp/data directory to support Unified Manager when you originally set up the Linux host and then installed Unified Manager, you can add disk space after installation by increasing disk space on the /opt/netapp/data directory.

Before you begin

You must have root user access to the Red Hat Enterprise Linux or CentOS Linux machine on which Unified Manager is installed.

About this task

We recommend that you back up the Unified Manager database before increasing the size of the data directory.

Steps

- 1. Log in as root user to the Linux machine on which you want to add disk space.
- 2. Stop the Unified Manager service and the associated MySQL software in the order shown: service ocieau stopservice ocie stopservice mysqld stop
- 3. Create a temporary backup folder (for example, /backup-data) with sufficient disk space to contain the data in the current /opt/netapp/data directory.
- 4. Copy the content and privilege configuration of the existing /opt/netapp/data directory to the backup data directory: cp -rp /opt/netapp/data/* /backup-data
- 5. If SE Linux is enabled:
 - a. Get the SE Linux type for folders on existing /opt/netapp/data folder:

```
se type= ls -Z /opt/netapp/data | awk '{print $4}'| awk -F: '{print $3}'|
```

The system returns a confirmation similar to the following:

```
echo $se_type
mysqld_db_t
```

- b. Run the choon command to set the SE Linux type for the backup directory: choon -R --type=mysqld db t /backup-data
- 6. Remove the contents of the /opt/netapp/data directory:

```
a. cd /opt/netapp/data
```

```
b. rm -rf *
```

7. Expand the size of the /opt/netapp/data directory to a minimum of 750 GB through LVM commands or by adding extra disks.



Mounting the /opt/netapp/data directory on an NFS export or CIFS share is not supported.

8. Confirm that the /opt/netapp/data directory owner (mysql) and group (root) are unchanged: ls -ltr / | grep opt/netapp/data

The system returns a confirmation similar to the following:

```
drwxr-xr-x. 17 mysql root 4096 Aug 28 13:08 data
```

9. If SE Linux is enabled, confirm that the context for the <code>/opt/netapp/data</code> directory is still set to <code>mysqld_db_t</code>: touch <code>/opt/netapp/data/abc``ls -Z /opt/netapp/data/abc</code>

The system returns a confirmation similar to the following:

```
-rw-r--r-. root root unconfined_u:object_r:mysqld_db_t:s0 /opt/netapp/data/abc
```

- 10. Copy the contents from backup-data, back to the expanded /opt/netapp/data directory: cp -rp /backup-data/* /opt/netapp/data/
- 11. Start the MySQL service: service mysqld start
- 12. After the MySQL service is started, start the ocie and ocieau services in the order shown: service ocie start``service ocieau start
- 13. After all of the services are started, delete the backup folder /backup-data: rm -rf /backup-data

Adding space to the data disk of the VMware virtual machine

If you need to increase the amount of space on the data disk for the Unified Manager

database, you can add capacity after installation by increasing disk space on disk 3.

Before you begin

- · You must have access to the vSphere Client.
- · The virtual machine must have no snapshots stored locally.
- · You must have the maintenance user credentials.

About this task

We recommend that you back up your virtual machine before increasing the size of virtual disks.

Steps

- 1. In the vSphere client, select the Unified Manager virtual machine, and then add more disk capacity to data disk 3. See the VMware documentation for details.
- 2. In the vSphere client, select the Unified Manager virtual machine, and then select the Console tab.
- 3. Click in the console window, and then log in to the maintenance console using your user name and password.
- 4. In the Main Menu, enter the number for the System Configuration option.
- 5. In the System Configuration Menu, enter the number for the Increase Data Disk Size option.

Adding space to the logical drive of the Microsoft Windows server

If you need to increase the amount of disk space for the Unified Manager database, you can add capacity to the logical drive on which Unified Manager is installed.

Before you begin

You must have Windows administrator privileges.

About this task

We recommend that you back up the Unified Manager database before adding disk space.

Steps

- 1. Log in as administrator to the Windows server on which you want to add disk space.
- 2. Follow the step that corresponds to method you want to use to add more space:

Option	Description
On a physical server, add capacity to the logical drive on which the Unified Manager server is installed.	Follow the steps in the Microsoft topic: Extend a Basic Volume
On a physical server, add a hard disk drive.	Follow the steps in the Microsoft topic: Adding Hard Disk Drives

Option	Description
On a virtual machine, increase the size of a disk partition.	Follow the steps in the VMware topic: Increasing the size of a disk partition

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

https://www.netapp.com/company/legal/copyright/

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

https://www.netapp.com/company/legal/trademarks/

Patents

A current list of NetApp owned patents can be found at:

https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf

Privacy policy

https://www.netapp.com/company/legal/privacy-policy/

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

Notice for OnCommand Unified Manager 9.5

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.