



OnCommand® Performance Manager 1.1

User Guide



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

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

Part number: 215-09107_A0
November 2014

Contents

Introduction to OnCommand Performance Manager	8
OnCommand Performance Manager features	8
OnCommand Performance Manager product documentation	9
Completing the setup wizard	10
Configuring your environment after deployment	11
Changing the Performance Manager host name	12
Configuring Performance Manager to send alert notifications	13
Customizing your environment	14
Enabling periodic AutoSupport	14
Sending an on-demand AutoSupport message	16
Configuring NTP settings	17
Editing network settings	17
Working with HTTPS security certificates	18
Protocol and port requirements	19
Page descriptions for system setup	20
AutoSupport dialog box	21
NTP Server dialog box	21
Configure Network Settings dialog box	22
Cluster concepts	24
What a cluster is	24
What a node in the cluster is	25
What an HA pair is	25
Connections and components of an HA pair	25
How HA pairs relate to the cluster	26
What logical storage is	29
What SVMs are	29
Why you use SVMs	31
What volumes are	32
Storage QoS	34
How the maximum throughput limit works	35
How throttling a workload can affect non-throttled workload requests from the same client	35

Controlling and monitoring I/O performance to FlexVol volumes by using Storage QoS	36
What network processing is	37
What data processing is	37
Preparing for the MetroCluster installation	37
Understanding the parts of the MetroCluster configuration	38
Replication of SVMs and switchover	41
Storage aggregates and disks	43
Aggregate states	43
Aggregate capacity states	44
Collecting data and monitoring workload performance	46
Types of workloads monitored by Performance Manager	47
Workload performance measurement values	48
What the expected range of performance is	50
How the expected range is used in performance analysis	51
How Performance Manager uses workload response time to identify performance issues	53
How cluster operations can affect workload response times	54
Performance monitoring of MetroCluster configurations	55
Volume behavior during switchover and switchback	56
What performance events are	58
Performance incident analysis and notification	59
How Performance Manager determines the performance impact for an incident	62
Why a cluster component can be in contention	63
Roles of workloads involved in a performance incident	65
Cluster configuration changes detected by Performance Manager	66
Navigating Performance Manager	68
Logging in to the GUI	69
Browser and platform requirements	70
How graphs of performance data work	70
Exporting data to a CSV file	72
Copying a link to a page	73
Printing a page	74
Analyzing workload performance	75
Determining whether a workload has a performance issue	75

Investigating a perceived slow response time for a workload	76
Identifying trends of I/O response time on cluster components	78
Analyzing the performance improvements achieved from moving a volume	79
How moving a FlexVol volume works	81
Volume Details page	82
Performance statistics displayed in the data breakdown charts	84
Analyzing performance incidents	87
Displaying information about a performance incident	87
Identifying victim workloads involved in a performance incident	88
Identifying bully workloads involved in a performance incident	89
Identifying shark workloads involved in a performance incident	91
Performance incident analysis for a MetroCluster configuration	91
Analyzing a performance incident on a cluster in a MetroCluster configuration	92
Checking the health of clusters in a MetroCluster configuration	95
Analyzing a performance incident for a remote cluster on a MetroCluster configuration	98
Responding to a performance incident caused by QoS policy group throttling	99
Responding to a performance incident caused by a disk failure	101
Responding to a performance incident caused by HA takeover	103
Page descriptions for analysis of performance incidents	105
Dashboard details for Quick Takes	105
Dashboard details for Incidents	107
Incident Details page	108
Managing data sources	117
Adding clusters	117
Requirements for adding a cluster to Performance Manager	118
How the discovery process works	118
Viewing the clusters list	119
Editing clusters	120
Removing clusters	120
Searching for storage objects	121
Page descriptions for data source management	122
Manage Data Sources page	122
Add Cluster dialog box	123
Edit Cluster dialog box	124

Managing users	126
What the maintenance user does	126
What RBAC is	126
What RBAC does	126
Authentication with Active Directory or OpenLDAP	127
Definitions of user types	128
Definitions of user roles in Performance Manager	128
Performance Manager user roles and capabilities	129
Adding users	129
Viewing users	131
Editing the user settings	131
Changing the local user password	132
Deleting users or groups	132
Page descriptions for user management	133
Manage Users page	133
Add User dialog box	134
Edit User dialog box	135
Managing user authentication	137
Authentication with Active Directory or OpenLDAP	137
Enabling remote authentication	138
Disabling nested groups from remote authentication	139
Setting up authentication services	139
Adding authentication servers	141
Editing authentication servers	143
Testing the configuration of authentication servers	143
Deleting authentication servers	144
Page descriptions for user authentication	145
Authentication dialog box	145
Managing security certificates	149
Viewing the HTTPS security certificate	149
Restarting the Performance Manager virtual machine	150
Generating an HTTPS security certificate	150
Downloading an HTTPS certificate signing request	151
Installing an HTTPS security certificate	152
Page descriptions for certificate management	153
HTTPS Certificate dialog box	153

Managing event notification	155
Configuring email settings	155
Configuring email alerts	156
Page descriptions for notification management	156
Email dialog box	156
Configure Email Alerts dialog box	157
Troubleshooting common issues	159
Unknown authentication error	159
Icons are misaligned in Internet Explorer	159
LDAP server slow to respond	160
Issue with adding LDAP using Other authentication services	160
Copyright information	162
Trademark information	163
How to send your comments	164
Index	165

Introduction to OnCommand Performance Manager

OnCommand Performance Manager provides performance monitoring and incident root-cause analysis for systems running clustered Data ONTAP software. It is the performance management part of OnCommand Unified Manager.

Performance Manager helps you identify workloads that are overusing cluster components and decreasing the performance of other workloads on the cluster. It alerts you to these performance events, called *incidents*, so that you can take corrective action and return performance back to normal operation. You can view and analyze incidents in the Performance Manager GUI or view them in the Unified Manager Dashboard.

Performance Manager monitors the performance of two types of workloads:

- User-defined workloads
These workloads consist of FlexVol volumes that you have created in your cluster.
- System-defined workloads
These workloads consist of internal system activity.

OnCommand Performance Manager features

OnCommand Performance Manager collects and analyzes performance statistics from a system running clustered Data ONTAP software. It uses a dynamic performance threshold to monitor the I/O response time of each volume on a cluster.

A high response time indicates that the volume is performing slower than normal. Slow response time also indicates that the performance has decreased for client applications that are using the volume. Performance Manager identifies the cluster component where the performance issue lies and provides a list of suggested actions you can take to try and address the performance issue.

Performance Manager can be deployed as a virtual appliance in VMware vSphere environments, or it can be installed on either physical or virtual servers running Red Hat Enterprise Linux.

OnCommand Performance Manager includes the following features:

- Monitors and analyzes workload performance statistics from a system running clustered Data ONTAP.
- Uses a dynamic performance threshold that analyzes the workload activity to identify and alert you to performance issues.
- Clearly identifies the cluster component that is in contention.

- Displays detailed graphs that plot workload activity over time, including I/O response time, the operations of user-defined and system-defined workloads, and cluster component usage.
- Identifies workloads that are overusing cluster components and the workloads whose performance is impacted by the increased activity.
- Identifies performance issues on clusters in a MetroCluster configuration.

OnCommand Performance Manager product documentation

OnCommand Performance Manager is accompanied by a set of guides that describe how to install and use the product.

OnCommand Performance Manager Installation and Administration Guide for VMware Virtual Appliances

Provides instructions for installing the Performance Manager appliance on a VMware ESXi server. This includes deploying and configuring the appliance and accessing the web-based interface.

OnCommand Performance Manager Installation and Setup Guide for Red Hat Enterprise Linux

Provides instructions for installing the Performance Manager appliance on a physical or virtual server running Red Hat Enterprise Linux. This includes deploying and configuring the appliance and accessing the web-based interface.

OnCommand Performance Manager User Guide

Provides an overview of Performance Manager, including reference information that explains the web-based interface, and instructions for monitoring, analyzing, and troubleshooting performance issues for workloads on a system running clustered Data ONTAP.

Completing the setup wizard

After the installation of Performance Manager is completed, you can access the GUI to complete the setup wizard. You use the setup wizard to configure email alerts, enable AutoSupport, change the password for the default user, and add clusters you want to monitor. After completing the setup wizard, you can access the other areas of the GUI.

Before you begin

- You are prepared to provide the setup wizard with the following information:
 - An IP address or URL at which to access the login window to the Performance Manager web UI
 - An assigned username for the maintenance user
 - A password for the maintenance user
- You can specify the following maintenance user information:
 - An email address at which the user can receive AutoSupport messages
 - The IP address of the associated SMTP mail server
- Clusters you want to add to Performance Manager meet configuration requirements.

Steps

1. Use an Internet browser to log in to the Performance Manager web UI, using the IP address or URL, the username for the maintenance user, and password that you were assigned at the end of installation.

The setup wizard is displayed, with four initial settings for you to configure.

2. Configure the following settings to complete the wizard:

Set Up Email

Specifies where email alerts are to be sent. You can identify an initial email recipient and an SMTP server to handle email communication. When Performance Manager is installed on Red Hat Enterprise Linux system, you can also specify a Network Time Protocol (NTP) server to synchronize Performance Manager server time with the time of the NTP server. After configuring the settings, you can click **Test** to confirm whether recipients can receive email alerts.

Set Up AutoSupport

Specifies whether AutoSupport is enabled to send information about your installation of Performance Manager to technical support. Support personnel can use this information to

stay current with the configuration and operation of your installation of Performance Manager, and to help you troubleshoot or manage the product.

Change Admin Credentials

Provides options for changing the password for the Administrator account. This is the password you assigned to the maintenance user account when you installed Performance Manager. You cannot change the Administrator account name.

When Performance Manager is installed on Red Hat Enterprise Linux system, you can change the password at a later time by logging into the Red Hat Enterprise Linux system as root and running the appropriate command.

Add Clusters

Enables you to add one or more clusters to monitor. You can enter the fully qualified name (FQDN) or IP address and access credentials for each system running clustered Data ONTAP. Each cluster must meet minimum configuration requirements.

Note: The first time you add a cluster, it can take up to 15 minutes for Performance Manager to fully discover it. Until the discovery process has completed, you cannot search for objects, such as volumes, on the cluster.

After you finish

If you choose not to immediately add clusters, you can configure additional options, such as alerts, and then add clusters for monitoring.

Related concepts

[What AutoSupport does](#) on page 15

[How the discovery process works](#) on page 118

[Browser and platform requirements](#) on page 70

Related tasks

[Adding clusters](#) on page 117

Configuring your environment after deployment

After you deploy the Performance Manager virtual appliance and complete the setup wizard, there are several configuration tasks that you might want to perform before you start monitoring your clusters, such as changing the host name, configuring alerts, and adding users.

Before you begin

- You must have deployed the virtual appliance and completed the initial setup of Performance Manager.

- You must be logged in as the OnCommand Administrator to complete all tasks in this workflow.

Choices

- [*Changing the Performance Manager host name*](#) on page 12

When you deployed Performance Manager, an SSL certificate was generated for HTTPS access. A host name was associated with the certificate, allowing you to use the host name to access the Performance Manager GUI. You might want to change this host name after deployment.

- [*Configuring Performance Manager to send alert notifications*](#) on page 13

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

- [*Adding users*](#) on page 129

You must manually add users to Performance Manager to create user accounts and control user access.

Changing the Performance Manager host name

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

Before you begin

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

About this task

Note: If Performance Manager is installed on Red Hat Enterprise Linux, you do not perform this workflow to change the host name, rather you change the host name using the Red Hat Linux command line.

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

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

If you access the Performance Manager GUI 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 that you do update the certificate, so that the host name in the certificate matches the actual host name.

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

Steps

1. [Edit the host name in Network Settings](#) on page 17

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

2. [Generate an HTTPS security certificate](#) on page 150

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

3. [View the HTTPS security certificate](#) on page 149

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

4. [Restart the Performance Manager virtual machine](#) on page 150

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

Configuring Performance Manager to send alert notifications

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

Before you begin

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

About this task

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

You can complete the following tasks to properly configure your environment to send alert notifications.

Steps

1. [Configure email settings](#) on page 155

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

2. [Enable remote authentication](#) on page 138

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

3. [Add authentication servers](#) on page 141

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

4. [Add users](#) on page 129

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. [Configure email alerts](#) on page 156

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

Customizing your environment

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

Enabling periodic AutoSupport

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

Before you begin

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

You must have defined the SMTP Server and Email settings in the Email dialog box. The system sends AutoSupport messages using this information.

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > AutoSupport**.

3. If you want to read about what periodic AutoSupport entails, click **View AutoSupport Description**.

The dialog box also displays the product System ID, which is the number that technical support uses to find your AutoSupport messages.

4. Select the **Enable Periodic AutoSupport** check box, and then click **Save and Close**.

Result

Automatic sending of AutoSupport messages is set up.

Note: AutoSupport messages are very large. It is possible that the email server will not send the message if it is larger than the allowable size. You might need to select a different SMTP host server if this problem occurs.

Related tasks

[Sending an on-demand AutoSupport message](#) on page 16

[Configuring email settings](#) on page 155

What AutoSupport does

With the help of the AutoSupport feature, Performance Manager sends information to technical support personnel to help with troubleshooting. AutoSupport (ASUP) messages are scanned for potential problems and are available to technical support personnel when they assist you in resolving issues.

When you generate the ASUP message from Performance Manager, the following configuration and analytical data is included in the ASUP message:

- Number of incidents that have a state of new or obsolete over the last seven days.
- Top three cluster components with the highest number of incidents over the last seven days.
- Configuration changes caused by HA takeover, policy group limit modifications, volume moves, or upgrade of the Data ONTAP software.
- Minimum, maximum, and average times for configuration and analytical data to be collected.
- Minimum, maximum, and average times for incident analysis to complete.
- Details about the virtual machine (VM), database, disk storage usage, and the number of errors and exceptions specific to Performance Manager.

ASUP or support messages that you generate through the maintenance console will not include the configuration and analytical data from Performance Manager.

Related tasks

[Enabling periodic AutoSupport](#) on page 14

[Sending an on-demand AutoSupport message](#) on page 16

Sending an on-demand AutoSupport message

You can choose to have Performance Manager send an on-demand message to technical support for assistance with troubleshooting issues. The AutoSupport message sent by Performance Manager contains diagnostic system information and detailed data about the Performance Manager server.

Before you begin

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

You have defined the SMTP Server and Email settings in the Email dialog box. The system sends AutoSupport messages using this information.

Steps

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

The dialog box also displays the product System ID, which is a unique ID for your Performance Manager instance that technical support uses to find your AutoSupport messages.

4. Click **Generate and Send AutoSupport**.

Result

An AutoSupport message is sent to technical support.

Note: AutoSupport messages are very large. It is possible that the email server will not send the message if it is larger than the allowable size. You may need to select a different SMTP host server if this problem occurs.

Related concepts

[What AutoSupport does](#) on page 15

Related tasks

[Enabling periodic AutoSupport](#) on page 14

[Configuring email settings](#) on page 155

Configuring NTP settings

You can use the NTP Server dialog box to specify the Network Time Protocol (NTP) server you want to use with Performance Manager. The Performance Manager server synchronizes its time with the time on the NTP server.

Before you begin

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

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > NTP Server**.
3. In the **NTP Server** dialog box, type the host name or FQDN, or the IP address of the NTP server.
Note: You can enter the name or address for more than one NTP server, but only the first entry is applied.
4. Click **Save and Close** to apply the setting.

Host names and FQDNs are resolved to IP addresses and stored as IP addresses.

Related tasks

[Completing the setup wizard](#) on page 10

Editing network settings

You might want to edit network settings if, for example, an IP address of a virtual machine (VM) changes, or if you switch from a DHCP to a static IP configuration.

Before you begin

Note: You cannot edit these settings when Performance Manager is installed on Red Hat Enterprise Linux.

- You will need one or more of the following networking items: host name or FQDN, IP address, DHCP, network mask, gateway, primary and secondary DNS addresses, and search domains.
- If you are changing your network settings from DHCP-enabled to static IP, you have verified the following items:
 - The IP address does not contain a duplicate address.
 - The gateway is reachable.

- The primary and secondary DNS addresses are ready and available to send and receive network traffic.
- You must have the OnCommand Administrator role or the Storage Administrator role.

Steps

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

Result

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

After you finish

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

Working with HTTPS security certificates

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

Before you begin

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

About this task

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

Choices

- [Generating an HTTPS security certificate](#) on page 150

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

- [Downloading an HTTPS certificate signing request](#) on page 151

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

- [Installing an HTTPS security certificate](#) on page 152

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

- [Viewing the HTTPS security certificate](#) on page 149

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

Protocol and port requirements

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

Connections to the Performance Manager server

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

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

Interface	Protocol	Port	Description
Performance Manager GUI	HTTP	80	Used to access the Performance Manager GUI; automatically redirects to the secure port 443.

Interface	Protocol	Port	Description
Performance Manager GUI and programs using APIs	HTTPS	443	Used to securely access the Performance Manager GUI 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.

Connections from the Performance Manager server

You must configure your firewall to open ports that enable communication between the Performance 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 Performance Manager server to connect to specific destinations.

The Performance 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.
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.
Mail server	SMTP	25	Used to send alert notification email.
Graphite server	TCP	2003	Used to send performance data.
NTP server	NTP	123/UDP	Used to synchronize the time on the Performance Manager server with an external NTP time server.

Page descriptions for system setup

You use the pages and dialog boxes in the GUI for configuring communication between the Performance Manager server and your network, and for enabling or disabling AutoSupport. When

enabled, AutoSupport routinely sends information about your Performance Manager instance to technical support.

AutoSupport dialog box

The AutoSupport dialog box enables you to view the AutoSupport description, enable periodic AutoSupport, or send an on-demand AutoSupport message. The dialog box also displays the product System ID, which is a unique ID for your Performance Manager instance that technical support uses to find your AutoSupport messages.

Information

You can perform the following operations:

View AutoSupport Description

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

Actions

You can perform the following operations:

Enable Periodic AutoSupport

Enables you to have specific, predefined messages to technical support periodically generated for issue diagnosis and resolution.

Generate and Send AutoSupport

Enables you to generate an on-demand message to send to technical support for any issues that have recently occurred.

Related concepts

[*What AutoSupport does*](#) on page 15

Related tasks

[*Enabling periodic AutoSupport*](#) on page 14

[*Sending an on-demand AutoSupport message*](#) on page 16

NTP Server dialog box

You can use the NTP Server dialog box to specify the NTP server that you want to use with Performance Manager. The Performance Manager server synchronizes its time with the time on the NTP server.

Note: You cannot edit these settings when Performance Manager is installed on Red Hat Enterprise Linux.

You can add, or change, the host name or IP address of the NTP server that you want to use.

Host Name or IP Address

Enables you to specify the host name or IP address of the NTP server.

Host names and FQDNs are resolved to IP addresses and stored as IP addresses.

Configure Network Settings dialog box

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

Note: You cannot edit these settings when Performance Manager is installed on Red Hat Enterprise Linux.

Host

The Host area provides the host name:

Host Name

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

Network

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

DHCP Enabled

Specifies whether DHCP is enabled.

Note: If DHCP is enabled, the system populates the IP address, network mask, and gateway fields with values from the network, and these fields appear dimmed. Also, you cannot change the values for the primary DNS address, the secondary DNS address, or the search domains.

IP Address

Specifies the IP address of the Performance Manager server.

Network Mask

Specifies the network mask.

Gateway

Specifies the IP address of the gateway.

Primary DNS Address

Specifies the IP address of the primary DNS server.

Secondary DNS Address

Specifies the IP address of the secondary DNS server.

Search Domains

Specifies the domain names (as comma-separated values) that are used by the DNS server to search for the host name.

Cluster concepts

To analyze workload performance activity and identify the cluster components involved in incidents, Performance Manager monitors and analyzes the logical and physical components of a cluster.

Related concepts

[Types of workloads monitored by Performance Manager](#) on page 47

What a cluster is

A cluster consists of one or more nodes grouped together as (HA pairs) to form a scalable cluster. Creating a cluster enables the nodes to pool their resources and distribute work across the cluster, while presenting administrators with a single entity to manage. Clustering also enables continuous service to end users if individual nodes go offline.

- The maximum number of nodes within a cluster depends on the platform model and licensed protocols.
- Each node in the cluster can view and manage the same volumes as any other node in the cluster. The total file-system namespace, which comprises all of the volumes and their resultant paths, spans the cluster.
- The nodes in a cluster communicate over a dedicated, physically isolated and secure Ethernet network.
The cluster logical interfaces (LIFs) on each node in the cluster must be on the same subnet.
- When new nodes are added to a cluster, there is no need to update clients to point to the new nodes.
The existence of the new nodes is transparent to the clients.
- If you have a two-node cluster (a single HA pair), you must configure cluster high availability (HA).
- You can create a cluster on a stand-alone node, called a single-node cluster.
This configuration does not require a cluster network, and enables you to use the cluster ports to serve data traffic. However, nondisruptive operations are not supported on single-node clusters.

Related concepts

[Cluster concepts](#) on page 24

What a node in the cluster is

A *node* is a controller in a cluster. It is connected to other nodes in the cluster over a private management cluster network. It is also connected to the disk shelves that provide physical storage for the Data ONTAP system or to third-party storage arrays that provide array LUNs for Data ONTAP use.

A *node Storage Virtual Machine (SVM)* represents a node in the cluster. The cluster setup process automatically creates a node SVM for each node in the cluster.

Related concepts

[Cluster concepts](#) on page 24

What an HA pair is

An HA pair is two storage systems (nodes) whose controllers are connected to each other directly. In this configuration, one node can take over its partner's storage to provide continued data service if the partner goes down.

You can configure the HA pair so that each node in the pair shares access to a common set of storage, subnets, and tape drives, or each node can own its own distinct set of storage.

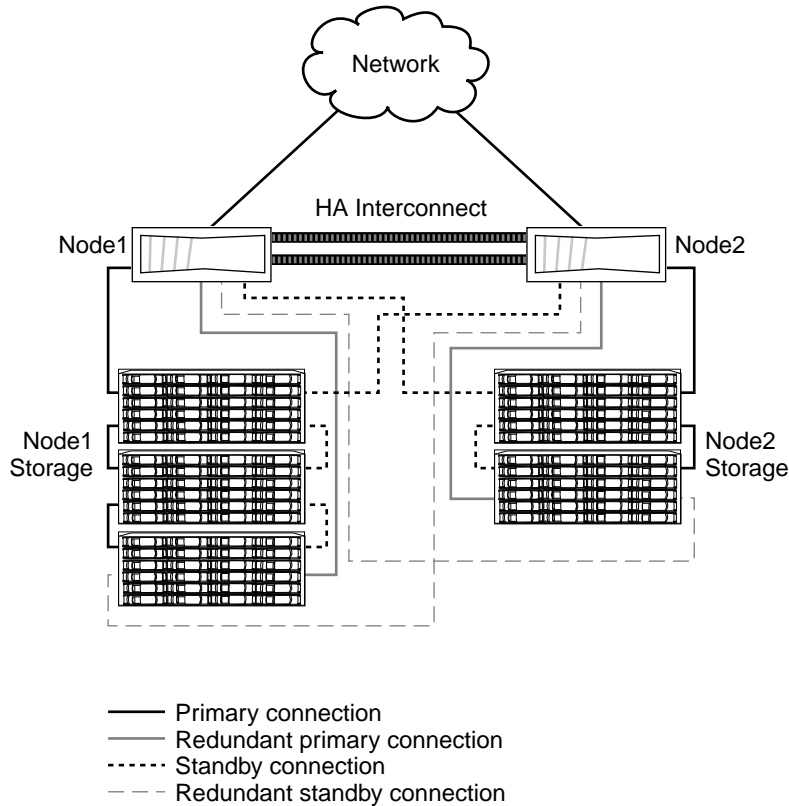
The controllers are connected to each other through an HA interconnect. This allows one node to serve data that resides on the disks of its failed partner node. Each node continually monitors its partner, mirroring the data for each other's nonvolatile memory (NVRAM or NVMEM). The interconnect is internal and requires no external cabling if both controllers are in the same chassis.

Takeover is the process in which a node takes over the storage of its partner. *Giveback* is the process in which that storage is returned to the partner. Both processes can be initiated manually or configured for automatic initiation.

Connections and components of an HA pair

Each node in an HA pair requires a network connection, an HA interconnect between the controllers, and connections to both its own disk shelves and its partner node's shelves.

The following diagram shows a standard HA pair with native DS4243 disk shelves and multipath HA:



How HA pairs relate to the cluster

HA pairs are components of the cluster, and both nodes in the HA pair are connected to other nodes in the cluster through the data and cluster networks. But only the nodes in the HA pair can take over each other's storage.

Although the controllers in an HA pair are connected to other controllers in the cluster through the cluster network, the HA interconnect and disk-shelf connections are found only between the node and its partner and their disk shelves or array LUNs.

The HA interconnect and each node's connections to the partner's storage provide physical support for high-availability functionality. The high-availability storage failover capability does not extend to other nodes in the cluster.


Note: Network failover does not rely on the HA interconnect and allows data network interfaces to failover to different nodes in the cluster outside the HA pair. Network failover is different than storage failover since it enables network resiliency across all nodes in the cluster.

Non-HA (or stand-alone) nodes are not supported in a cluster containing two or more nodes. Although single-node clusters are supported, joining two separate single-node clusters to create one

cluster is not supported, unless you wipe clean one of the single-node clusters and join it to the other to create a two-node cluster that consists of an HA pair.

Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators

The following diagram shows two HA pairs. The multipath HA storage connections between the nodes and their storage are shown for each HA pair. For simplicity, only the primary connections to the data and cluster networks are shown.



Key to storage connections

- Primary connection
- Redundant primary connection
- - - Standby connection
- - - Redundant standby connection

Possible storage failover scenarios in this cluster are as follows:

- Node1 fails and Node2 takes over Node1's storage.
- Node2 fails and Node1 takes over Node2's storage.
- Node3 fails and Node4 takes over Node3's storage.
- Node4 fails and Node3 takes over Node4's storage.

If Node1 *and* Node2 both fail, the storage owned by Node1 and Node2 becomes unavailable to the data network. Although Node3 and Node4 are clustered with Node1 and Node2, they do not have direct connections to Node1 and Node2's storage and cannot take over their storage.

What logical storage is

Logical storage refers to the storage resources provided by Data ONTAP that are not tied to a physical resource.

Logical storage resources are associated with a Storage Virtual Machine (SVM, formerly known as Vserver), and they exist independently of any specific physical storage resource such as a disk, array LUN, or aggregate. Logical storage resources include volumes of all types and qtrees, as well as the capabilities and configurations you can use with these resources, such as Snapshot copies, deduplication, compression, and quotas.

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


What SVMs are

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

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

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

SVM with FlexVol volumes




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

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

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

SVM with Infinite Volume



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

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

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

Why you use SVMs

Storage Virtual Machines (SVMs, formerly known as Vservers) provide data access to clients regardless of the physical storage or controller, similar to any storage system. SVMs provide benefits such as nondisruptive operations, scalability, security, and unified storage.

SVMs provide the following benefits:

- **Multi-tenancy**
SVM is the fundamental unit of secure multi-tenancy, which enables partitioning of the storage infrastructure so that it appears as multiple independent storage systems. These partitions isolate the data and management.
- **Nondisruptive operations**
SVMs can operate continuously and nondisruptively for as long as they are needed. SVMs help clusters to operate continuously during software and hardware upgrades, addition and removal of nodes, and all administrative operations.
- **Scalability**
SVMs meet on-demand data throughput and the other storage requirements.
- **Security**
Each SVM appears as a single independent server, which enables multiple SVMs to coexist in a cluster while ensuring no data flows among them.
- **Unified storage**
SVMs can serve data concurrently through multiple data access protocols. SVMs provide file-level data access through NAS protocols, such as CIFS and NFS, and block-level data access through SAN protocols, such as iSCSI and FC (FCoE included). SVMs can serve data to SAN and NAS clients independently at the same time.

Note: SVMs with Infinite Volume can serve data only through NFS and CIFS protocols.

- **Easy management of large datasets**
With SVMs with Infinite Volume, management of large and unstructured data is easier because the SVM administrator can manage one data container instead of many.

What volumes are

Volumes are data containers that enable you to partition and manage your data. Understanding the types of volumes and their associated capabilities enables you to design your storage architecture for maximum storage efficiency and ease of administration.

Volumes are the highest-level logical storage object. Unlike aggregates, which are composed of physical storage resources, volumes are completely logical objects.

System Manager supports two types of volumes, traditional and flexible. However, you can create only flexible volumes (FlexVol volumes) by using System Manager.

How volumes work

Volumes are data containers that enable you to partition and manage your data. Understanding the types of volumes and their associated capabilities enables you to design your storage architecture for maximum storage efficiency and ease of administration.

Volumes are the highest-level logical storage object. Unlike aggregates, which are composed of physical storage resources, volumes are completely logical objects.

Data ONTAP provides two types of volumes: FlexVol volumes and Infinite Volumes. There are also volume variations, such as FlexClone volumes, data protection mirrors, and load-sharing mirrors. Not all volume variations are supported for both types of volumes. Data ONTAP efficiency capabilities, compression and deduplication, are supported for both types of volumes.

Volumes contain file systems in a NAS environment, and LUNs in a SAN environment.

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

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

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

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

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

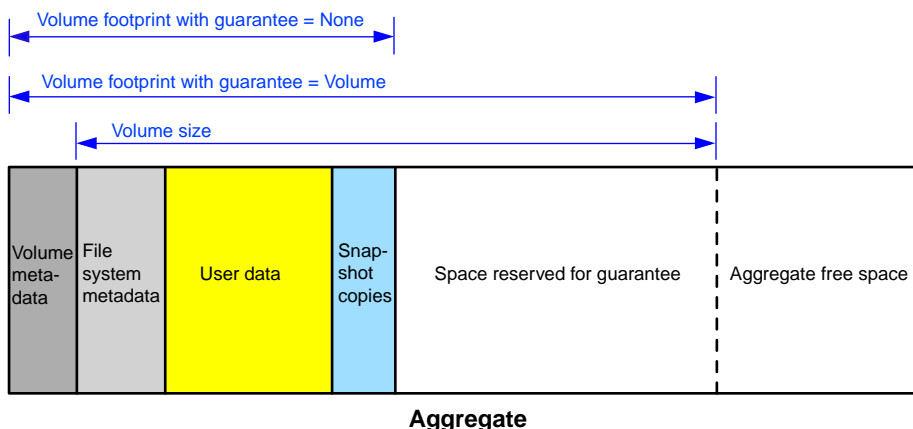
For more information about data protection mirrors, see the *Clustered Data ONTAP Data Protection Guide*.

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

What the volume footprint is

A volume footprint is the amount of space a volume is using within the aggregate. Understanding what is included in the volume footprint helps you understand the space requirements for the volume.

The volume footprint consists of the space used by user data and metadata, including metadata that resides in the aggregate rather than within the volume itself. For this reason it can be larger than the volume size, as shown in the following diagram:



Volume states

Volumes can be in one of four states—online, offline, restricted, or mixed.

On the Volume Details page, the volume status is displayed in parentheses at the top of the page next to the volume name.

The following table displays the possible states for volumes.

State	Description
online	Read and write access to this volume is allowed.
offline	No access to the volume is allowed.
restricted	Some operations, such as parity reconstruction, are allowed, but data access is not allowed.
mixed	The constituents of an Infinite Volume are not all in the same state.

Considerations for using thin provisioning with FlexVol volumes

Using thin provisioning, you can configure your volumes so that they appear to provide more storage than they have available, provided that the storage that is actually being used does not exceed the available storage.

To use thin provisioning with FlexVol volumes, you create the volume with a guarantee of **none**. With a guarantee of **none**, the volume size is not limited by the aggregate size. In fact, each volume could, if required, be larger than the containing aggregate. The storage provided by the aggregate is used up only as data is written to the LUN or file.

If the volumes associated with an aggregate show more storage as available than the physical resources available to that aggregate, the aggregate is *overcommitted*. When an aggregate is overcommitted, it is possible for writes to LUNs or files in volumes contained by that aggregate to fail if there is not sufficient free space available to accommodate the write.

If you have overcommitted your aggregate, you must monitor your available space and add storage to the aggregate as needed to avoid write errors due to insufficient space.

Aggregates can provide storage to FlexVol volumes associated with more than one Storage Virtual Machine (SVM). When sharing aggregates for thin-provisioned volumes in a multi-tenancy environment, be aware that one tenant's aggregate space availability can be adversely affected by the growth of another tenant's volumes.

Storage QoS

Data ONTAP 8.2 introduces Storage QoS, which can help you manage the risks that accompany meeting performance objectives for workloads. You can use Storage QoS to limit the throughput to workloads to a Storage Virtual Machine (SVM), or to groups of volumes or LUNs within an SVM, and to monitor IOPS and MBps performance.

You can reactively limit workload performance to ensure fair resource usage. If you have a cloud infrastructure, you might proactively limit workloads, as defined by their service levels.

For example, you can prevent runaway workloads from impacting other workloads in a shared storage infrastructure by applying a throughput limit to the runaway workload. In a service-provider environment, you can proactively set throughput limits at an SVM level, where an SVM maps to a tenant. This throughput limit ensures a consistent performance for each tenant as you add more tenants to the shared storage infrastructure and prevents tenants from affecting each other's workload performance.

Storage QoS is supported on clusters that have up to eight nodes.

For information about how to use Storage QoS, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

How the maximum throughput limit works

You can specify one service-level objective for a Storage QoS policy group: a maximum throughput limit. A maximum throughput limit, which you define in terms of IOPS or MBps, specifies the throughput that the workloads in the policy group cannot collectively exceed.

When you specify a maximum throughput for a policy group, Storage QoS controls client operations to ensure that the combined throughput for all workloads in the policy group does not exceed the specified maximum throughput.

For example, assume that you create the policy group “untested_apps” and specify a maximum throughput of 300 MBps. You assign three volumes to the policy group. The combined throughput to those three volumes cannot exceed 300 MBps.

Note: The combined throughput to the workloads in a policy group might exceed the specified limit by up to 10%. A deviation might occur if you have a workload that experiences rapid changes in throughput (sometimes called a *bursty workload*).

Note the following about specifying a maximum throughput:

- A throughput limit applies to all clients that access a storage object.
- Do not set the limit too low, because you might underutilize the cluster.
- Consider the minimum amount of throughput that you want to reserve for workloads that do not have limits.

For example, you can ensure that your critical workloads get the throughput that they need by limiting noncritical workloads.

- You might want to provide room for growth.
For example, if you see an average utilization of 500 IOPS, you might specify a limit of 1,000 IOPS.

How throttling a workload can affect non-throttled workload requests from the same client

In some situations, throttling a workload (I/O to a storage object) can affect the performance of non-throttled workloads if the I/O requests are sent from the same client.

If a client sends I/O requests to multiple storage objects and some of those storage objects belong to Storage QoS policy groups, performance to the storage objects that do not belong to policy groups might be degraded. Performance is affected because resources on the client, such as buffers and outstanding requests, are shared.

For example, this might affect a configuration that has multiple applications or virtual machines running on the same host.

This behavior is likely to occur if you set a low maximum throughput limit and there are a high number of I/O requests from the client.

If this occurs, you can increase the maximum throughput limit or separate the applications so they do not contend for client resources.

Controlling and monitoring I/O performance to FlexVol volumes by using Storage QoS

You can control input/output (I/O) performance to FlexVol volumes by assigning volumes to Storage QoS policy groups. You might control I/O performance to ensure that workloads achieve specific performance objectives or to throttle a workload that negatively impacts other workloads.

About this task

Policy groups enforce a maximum throughput limit (for example, 100 MB/s). You can create a policy group without specifying a maximum throughput, which enables you to monitor performance before you control the workload.

You can also assign Storage Virtual Machines (SVMs) with FlexVol volumes, LUNs, and files to policy groups.

Note the following requirements about assigning a volume to a policy group:

- The volume must be contained by the SVM to which the policy group belongs. You specify the SVM when you create the policy group.
- If you assign a volume to a policy group, then you cannot assign the volume's containing SVM or any child LUNs or files to a policy group.

Note: Storage QoS is supported on clusters that have up to eight nodes.

For more information about how to use Storage QoS, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Use the `qos policy-group create` command to create a policy group.
2. Use the `volume create` command or the `volume modify` command with the `-qos-policy-group` parameter to assign a volume to a policy group.
3. Use the `qos statistics` commands to view performance data.
4. If necessary, use the `qos policy-group modify` command to adjust the policy group's maximum throughput limit.

What network processing is

Network processing is a software component in the cluster that handles read and write requests between client applications and the network protocols on the cluster.

The network processing operations communicate with the data processing software component to locate the storage aggregate that will fulfill the request. Once the request is fulfilled, the data processing component communicates with the network processing component to transfer the information to the requesting client.

What data processing is

Data processing is a software component in the cluster that handles read and write requests to the target storage aggregate.

The data processing component receives read and write requests from the network processing software component. It locates the storage aggregate that can fulfill the requests, and communicates with the network processing component to transfer information to the requesting client applications.

Preparing for the MetroCluster installation

As you prepare for the MetroCluster installation, you should understand the MetroCluster hardware architecture and required components. If you are familiar with MetroCluster configurations in a 7-mode environment, you should understand the key MetroCluster differences you find in a clustered Data ONTAP environment.

Related concepts

[*Performance monitoring of MetroCluster configurations*](#) on page 55

Related tasks

[*Analyzing a performance incident on a cluster in a MetroCluster configuration*](#) on page 92

[*Checking the health of clusters in a MetroCluster configuration*](#) on page 95

[*Analyzing a performance incident for a remote cluster on a MetroCluster configuration*](#) on page 98

[*Identifying victim workloads involved in a performance incident*](#) on page 88

[*Identifying bully workloads involved in a performance incident*](#) on page 89

[*Identifying shark workloads involved in a performance incident*](#) on page 91

Related references

[Performance incident analysis and notification](#) on page 59


Understanding the parts of the MetroCluster configuration

The MetroCluster configuration consists of a *disaster recovery (DR) group* that includes two HA pairs, each in a separate cluster at physically separated sites. FC switches and long distance inter-switch links provide a backbone connection between the clusters. The clusters are also in a peering relationship, with each cluster's configuration information mirrored to the partner

The MetroCluster configuration includes the following key hardware elements:

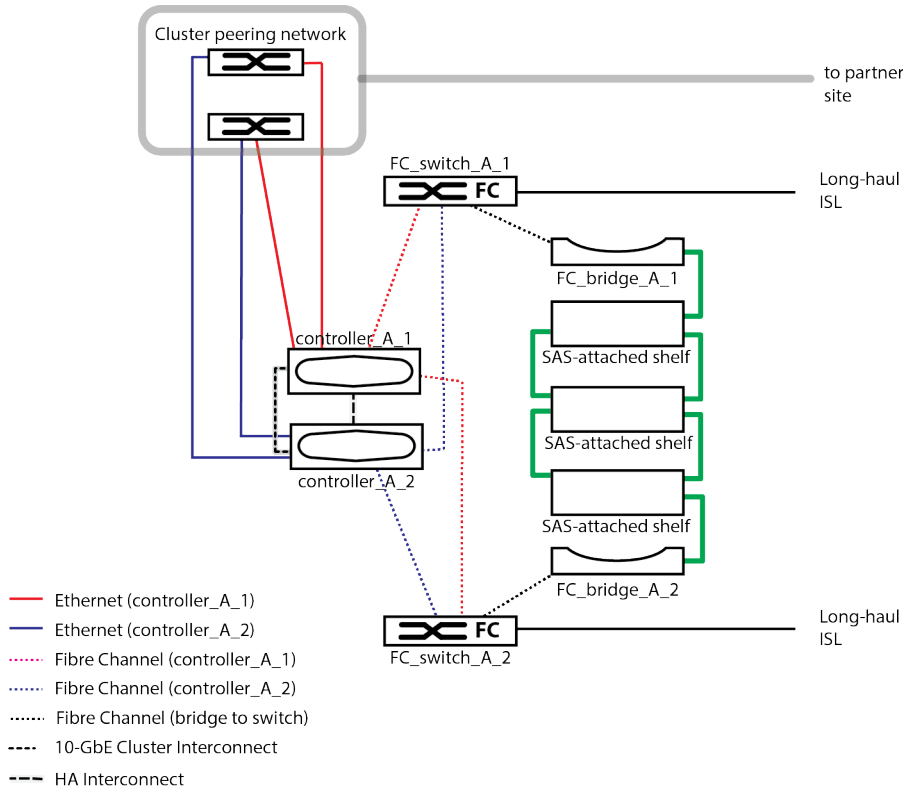
- **Storage controllers**
The storage controllers are not connected directly to the storage but connect to two redundant FC switch fabrics.
- **FC-to-SAS bridges**
The FC-to-SAS bridges connect the SAS storage stacks to the FC switches, providing bridging between the two protocols.
- **FC switches**
The FC switches provide the long haul backbone ISL between the two sites. The 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 Storage Virtual Machine (SVM) mirroring.

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 clusters, one at each geographically separated site.
- cluster_A is located at one MetroCluster site.
- cluster_B is located at the second MetroCluster site.
- Each site has one stack of SAS storage.
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 not shown.

The following illustration shows a more detailed view of the connectivity in a single MetroCluster cluster (both clusters have the same configuration):



The configuration includes the following connections:

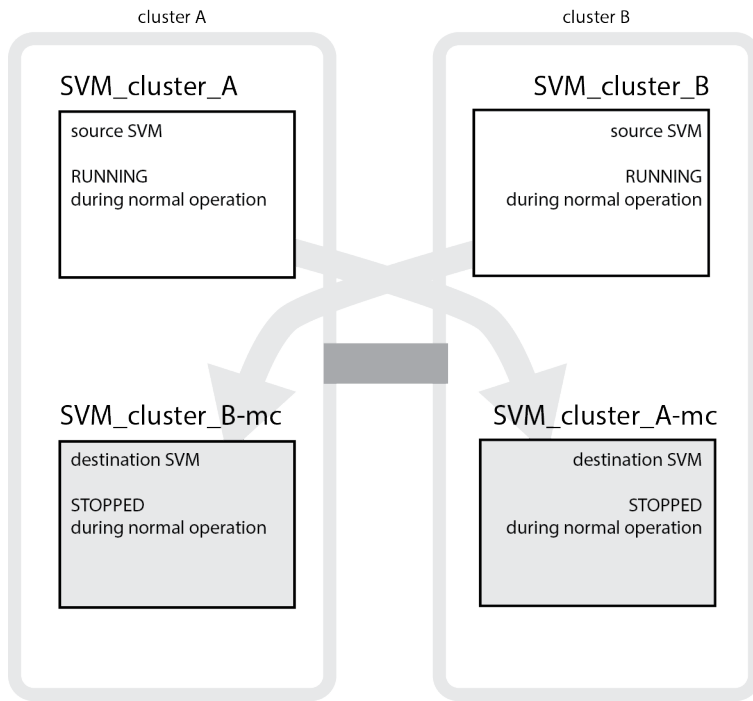
- FC connections from each controller's HBAs and FC-VI adapters to each of the FC switches.
- An FC connection from each FC-to-SAS bridge to an FC switch.
- SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge.
- An HA interconnect between each controller in the local HA pair.
If the controllers support a single-chassis HA pair, the HA interconnect is internal, occurring through the backplane, meaning an external interconnect is not required.
- Ethernet connections from the controllers to the customer-provided network used for cluster peering.
SVM configuration is replicated over the cluster peering network.
- A cluster interconnect between each controller in the local HA pair.
If the controllers are configured as a switched cluster, each controller would connect to two cluster interconnect switches.

Replication of SVMs and switchover


SVM mirroring provides redundant data server configuration and mirroring of data volumes that belong to the SVM. If a switchover occurs, the source SVM is brought down and the destination SVM, located on the surviving cluster, becomes active.

The following example shows the SVMs for a MetroCluster configuration, where vs1 is a SVM on the source site and vs-1mc is a sync-destination on the disaster recovery site (MetroCluster appends -mc to the name of the destination SVMs):

- vs1 serves data on cluster A.
It is a sync-source SVM that replicates the SVM configuration (LIFs, protocols, and services) and data in volumes belonging to the SVM. The configuration and data are replicated to vs1-mc, a sync-destination SVM located on cluster B.
- vs2 serves data on cluster B.
It is a sync-source SVM that replicates configuration and data to vs2-mc located on cluster A.
- vs2-mc is a sync-destination that is stopped during normal, healthy operation of the MetroCluster configuration.
In a successful switchover from cluster B to cluster A, vs2 is stopped and vs2-mc is activated and begins serving data from cluster A.
- vs1-mc is a sync-destination that is stopped during normal, healthy operation of the MetroCluster configuration.
In a successful switchover from cluster A to cluster B, vs1 is stopped and vs1-mc is activated and begins serving data from cluster B.



If a switchover occurs, the remote plex on the surviving cluster comes online and the secondary SVM begins serving the data.



Storage aggregates and disks

A storage aggregate is a logical grouping of physical storage that is protected by RAID technology. When you create a storage aggregate, you specify the number of disks in the underlying RAID set and the level of RAID protection for the set (RAID-4 or RAID-DP).

A storage aggregate contains a defined amount of physical storage that can be expanded dynamically at any time. For example, if a storage aggregate is made up of three disks, an additional disk can be easily added to the storage aggregate, and the logical size of the storage aggregate can be increased accordingly, without disrupting any user or process currently using that storage aggregate. Storage aggregates can include hot spare disks that hold data if one of the other disks fails.

Aggregate states


The state of an aggregate indicates its availability or whether it is involved in a specific process.



State	Description
Offline	Read or write access is not allowed.

State	Description
Restricted	Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.
Online	Read and write access to volumes hosted on this aggregate is allowed.
Creating	The aggregate is being created.
Destroying	The aggregate is being destroyed.
Failed	The aggregate cannot be brought online.
Frozen	The aggregate is (temporarily) not serving requests.
Inconsistent	The aggregate has been marked corrupted; contact technical support.
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.
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 OnCommand application server.

Aggregate capacity states

The storage capacity of an aggregate can be in 1 of 3 states: normal, warning, or error. The states are based on pre-defined capacity thresholds.

State	Description
 Normal	Used capacity is under the Warning and Error thresholds.

State	Description
 Warning	Used capacity is above the Warning threshold of 85% of the total aggregate capacity.
 Error	Used capacity is above the Error threshold of 95% of the total aggregate capacity.

Collecting data and monitoring workload performance

Performance Manager collects and analyzes workload activity every 5 minutes to identify performance incidents, and it detects configuration changes every 15 minutes. It retains a maximum of 90 days of historical performance and event data, and it uses this data to forecast the expected range for all monitored workloads.

Performance Manager must collect a minimum of 3 hours, or 36 data samples, 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 Volume Details page. While this activity is being collected, the expected range does not display all changes occurring from workload activity. After collecting 3 hours of activity, Performance 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 Performance Manager is monitoring a volume, if more than 24 hours have passed since the last data collection, the charts on the Volume Details page will not display the expected range for that volume. Incidents detected prior to the last collection are still available.

Note: Daylight savings time (DST) changes the system time, which alters the expected range of performance statistics for monitored workloads. Performance Manager immediately begins to correct the expected range, which takes approximately 15 days to complete. During this time you can continue to use Performance Manager, but, since Performance Manager uses the expected range to detect incidents, some incidents might not be accurate. Incidents detected prior to the time change are not affected. Manually changing the time on a Data ONTAP cluster, or a Performance Manager server, to an earlier time will also affect the incident analysis results.

Related concepts

[*How the discovery process works*](#) on page 118

[*What the expected range of performance is*](#) on page 50

Related tasks

[*Completing the setup wizard*](#) on page 10

Related references

[*Workload performance measurement values*](#) on page 48

Types of workloads monitored by Performance Manager

You can use Performance 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 is a user-defined workload.

Note: Performance 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.

Performance Manager requires that the volumes you want to monitor be in a QoS policy group:

- When using clustered Data ONTAP 8.3, a policy group is assigned to all volumes, either by the administrator or by Data ONTAP.
- When using clustered Data ONTAP 8.2.x, a policy group is assigned to all volumes, either by the administrator or by Performance Manager when the cluster is added to the UI. When Performance Manager analyzes the cluster for configuration changes every 15 minutes, it adds any new volumes not in a policy group to the default policy group.

Note: With clustered Data ONTAP 8.2.x, if an SVM, LUN, or File storage object is in a policy group, Performance Manager cannot monitor the volumes contained in that object and the overall analysis is impacted. You must remove the storage object from the policy group to correct this issue.

If one or more of the following is true for a workload, it cannot be monitored by Performance Manager:

- It is a data protection copy in read-only mode.
- 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.

For more information about workloads in storage QoS, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Related concepts

[Roles of workloads involved in a performance incident](#) on page 65

[Cluster concepts](#) on page 24

[Why a cluster component can be in contention](#) on page 63

Related references


[Performance incident analysis and notification](#) on page 59

[Incident Details page](#) on page 108

Workload performance measurement values

Performance 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 incident 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	<p>The percentage of the QoS limit used by the workloads in the policy group.</p> <p>Note: If Performance 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%.</p>
Actual	The measured performance value at a specific time for a given workload.
Deviation	<p>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.</p> <p>Note: 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.</p>
Expected	The expected values are based on the analysis of historical performance data for a given workload. Performance Manager analyzes these statistical values to determine the expected range of values.

Measurement	Description
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 response time, the upper values form the performance threshold. When the actual value crosses the performance threshold, Performance Manager triggers a performance incident alert.
Peak	The maximum value measured over a period of time.
Peak Deviation	The maximum deviation value measured over a period of time.
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.

Related concepts

[Collecting data and monitoring workload performance](#) on page 46

[What the expected range of performance is](#) on page 50

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 response time, the upper values form the performance threshold. When the actual value crosses the performance threshold, Performance Manager triggers a performance incident 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

Performance Manager must collect a minimum of 3 hours, or 36 data samples, 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 hours of activity, Performance 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.

Note: Daylight savings time (DST) changes the system time, which alters the expected range of performance statistics for monitored workloads. Performance Manager immediately begins to correct the expected range, which takes approximately 15 days to complete. During this time you can continue to use Performance Manager, but, since Performance Manager uses the expected range to detect incidents, some incidents might not be accurate. Incidents detected prior to the time change are not affected. Manually changing the time on a Data ONTAP cluster, or a Performance Manager server, to an earlier time will also affect the incident analysis results.

Related concepts

[Collecting data and monitoring workload performance](#) on page 46

Related tasks

[Determining whether a workload has a performance issue](#) on page 75

Related references

[How the expected range is used in performance analysis](#) on page 51

[Performance incident analysis and notification](#) on page 59

[Workload performance measurement values](#) on page 48

How the expected range is used in performance analysis

Performance Manager uses the expected range to represent the typical I/O response time and operations activity for your monitored workloads. It alerts you when the actual response time for a workload is above the upper bounds of the expected range, which triggers a performance incident, 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, Performance 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. Performance Manager uses the baseline to determine when the actual response time 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 response time for a workload exceeds the performance threshold, due to contention on a cluster component, the response time 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.

Performance Manager analyzes the threshold crossing event and determines whether the activity is a performance incident. If the high workload activity remains consistent for a long period of time, such as several hours, Performance 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 response time does not have a high rate of change over time. To minimize the number of incident alerts, during analysis of performance incidents, Performance Manager triggers an incident only for low-activity volumes whose operations and response times are much higher than expected.



In this example, the response time 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 response time, 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 response time returns within the expected range. If the response time remains at or above 10 ms/op for a long period of time, you might need to take corrective action to resolve the incident.

Related concepts

[What the expected range of performance is](#) on page 50

[Collecting data and monitoring workload performance](#) on page 46

Related tasks

[Determining whether a workload has a performance issue](#) on page 75

Related references

[Performance incident analysis and notification](#) on page 59

Workload performance measurement values on page 48

How Performance Manager uses workload response time to identify performance issues


The workload response time is the time it takes for a volume on a cluster to respond to I/O requests from client applications. Performance Manager uses the response time to detect and alert you to performance incidents.

A high response time means that requests from applications to a volume on a cluster are taking longer than usual. The cause of the high response time could be on the cluster itself, due to contention on one or more cluster components. High response time 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.

Note: Performance 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 response time. If the actual response time exceeds the performance threshold of the expected range, Performance Manager analyzes the event to determine if it is a performance incident that you might need to resolve. The response time is measured in milliseconds per operation (ms/op).

On the Volume Details page, you can view an analysis of the response time statistics to see how the activity of individual processes, such as read and write requests, compares to the overall response time statistics. The comparison helps you determine which operations have the highest activity or whether specific operations have abnormal activity that is impacting the response time for a volume. When analyzing performance incidents, you can use the response time statistics to determine whether an incident was caused by an issue on the cluster. You can also identify the specific workload activities or cluster components that are involved in the incident.



This example shows the Response Time chart on the Volume Details page. The actual response time activity is a blue line and the expected range is gray.

Note: There can be gaps in the blue line if Performance Manager was unable to gather data. This can occur because the cluster or volume was unreachable, Performance Manager was turned off during that time, or the collection was taking longer than the 5 minute collection period.

Related concepts

[What the expected range of performance is](#) on page 50

[How cluster operations can affect workload response times](#) on page 54

[How graphs of performance data work](#) on page 70

Related tasks

[Determining whether a workload has a performance issue](#) on page 75

[Investigating a perceived slow response time for a workload](#) on page 76

Related references

[Performance incident analysis and notification](#) on page 59

How cluster operations can affect workload response times

Operations represent the activity of all user-defined and system-defined workloads on a cluster. The operations statistics help you determine whether cluster processes, such as making backups or running deduplication, are impacting workload response times or might have caused, or contributed to, a performance incident.

When analyzing performance incidents, you can use the operations statistics to determine if a performance incident 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 incident. Operations are measured in operations per second (ops/sec).



This example shows the Operations chart on the Volume Details page. The actual operations statistics is a blue line and the expected range of operations statistics is gray.

Note: In some cases where a cluster is overloaded, Performance Manager may display the message Data collection is taking too long on Cluster <cluster_name>.

This means that not enough statistics have been collected for Performance Manager to analyze. You need to reduce the resources the cluster is using so that statistics can be collected.

Related concepts

[What the expected range of performance is](#) on page 50

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

[How graphs of performance data work](#) on page 70

Related tasks

[Investigating a perceived slow response time for a workload](#) on page 76

Performance monitoring of MetroCluster configurations

Performance 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, Performance Manager triggers performance incidents 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. Performance 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, Performance Manager does not analyze read throughput.

Performance 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.

Related concepts

[Preparing for the MetroCluster installation](#) on page 37

[Roles of workloads involved in a performance incident](#) on page 65

Related tasks

[Analyzing a performance incident on a cluster in a MetroCluster configuration](#) on page 92

[Checking the health of clusters in a MetroCluster configuration](#) on page 95

[Analyzing a performance incident for a remote cluster on a MetroCluster configuration](#) on page 98

[Identifying victim workloads involved in a performance incident](#) on page 88

[Identifying bully workloads involved in a performance incident](#) on page 89

[Identifying shark workloads involved in a performance incident](#) on page 91

Related references


[Performance incident analysis and notification](#) on page 59

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. Performance Manager monitors only those volumes that are active and running.

Since volumes are moved from one cluster to another, it is recommended that you monitor both clusters. A single instance of Performance Manager can monitor both clusters in a MetroCluster configuration, but sometimes the distance between the two locations necessitates using two Performance Manager instances to monitor both clusters. The following figure shows a single instance of Performance Manager:


Normal operation



Cluster B fails --- switchover to Cluster A



Cluster B is repaired --- switchback to Cluster B



□ = active and monitored by OPM ■ = inactive and not monitored by OPM

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: Vol1_p and Vol2_p.
- Cluster B has two active volumes: Vol3_p and Vol4_p.
- Cluster A has two inactive volumes: Vol3_b and Vol4_b.
- Cluster B has two inactive volumes: Vol1_b and Vol2_b.

Information pertaining to each of the active volumes (statistics, incidents, and so on) is collected by Performance Manager. Vol1_p and Vol2_p statistics are collected by Cluster A, and Vol3_p and Vol4_p 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: Vol1_p, Vol2_p, Vol3_b, and Vol4_b.
- Cluster B has four inactive volumes: Vol3_p, Vol4_p, Vol1_b, and Vol2_b.

As during normal operation, information pertaining to each of the active volumes is collected by Performance Manager. But in this case, Vol1_p and Vol2_p statistics are collected by Cluster A, and Vol3_b and Vol4_b statistics are also collected by Cluster A.

Note that Vol3_p and Vol3_b are not the same volumes, because they are on different clusters. The information in Performance Manager for Vol3_p is not the same as Vol3_b:

- During switchover to Cluster A, Vol3_p statistics and incidents are not visible.
- On the very first switchover, Vol3_b looks like a new volume with no historical information.

When Cluster B is repaired and a switchback is performed, Vol3_p is active again on Cluster B, with the historical statistics and a gap of statistics for the period during the switchover. Vol3_b is not viewable from Cluster A until another switchover occurs:



Note:

- MetroCluster volumes that are inactive, for example, Vol3_b on Cluster A after switchback, will be identified with the message “This volume was deleted”. The volume is not actually deleted, but it is not currently being monitored by Performance Manager because it is not the active volume.
- If a single Performance Manager is monitoring both clusters in a MetroCluster configuration, volume search will return information for whichever volume is active at that time. For example, a search for “Vol3” would return statistics and incidents for Vol3_b on Cluster A if a switchover has occurred and Vol3 has become active on Cluster A.

What performance events are

Performance events are incidents or configuration changes related to workload performance on a cluster. They help you identify workloads with slow response times and cluster configuration changes that might have caused or contributed to the slow response times.

When Performance 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.

Types of performance events

Performance Manager can detect the following types of performance events:

Incidents

Performance issues that are the result of failures or errors in an IT infrastructure. The cause of incidents might be a simple issue that corrects itself over a period of time or can be addressed with a repair or configuration change. In Performance Manager, an incident indicates workloads on a system running clustered Data ONTAP that are slow due to other workloads with high usage of shared cluster components. If the response time of the impacted workloads does not return to normal levels, you might need to take corrective action to resolve the incident.

Changes

The addition, modification, or removal of storage objects in an IT infrastructure. In Performance Manager, a change is an addition, modification, or deletion of a logical or physical storage object in a system running clustered Data ONTAP. A change that impacts the performance of one or more workloads might be a contributor to an incident.

Related concepts

[*What the expected range of performance is*](#) on page 50

[*Cluster configuration changes detected by Performance Manager*](#) on page 66

[*Roles of workloads involved in a performance incident*](#) on page 65

[*Types of workloads monitored by Performance Manager*](#) on page 47

Related tasks

[Displaying information about a performance incident](#) on page 87

Related references

[Performance incident analysis and notification](#) on page 59

Performance incident analysis and notification

Performance incidents are events that notify you about I/O performance issues on a volume workload caused by contention on a cluster component. Performance Manager analyzes the incident to identify all workloads involved, the component in contention, and whether the incident is still an issue that you might need to resolve.

Performance Manager monitors the I/O response time and operations for volumes on a cluster. When other workloads overuse a cluster component, for example, the component is in contention and is unable to perform at an optimal level to meet workload demands. The performance of other workloads that are using the same component might be impacted, causing their response times to increase. If the response time crosses the performance threshold, Performance Manager triggers a performance incident and sends an email alert to notify you.

Incident analysis

Performance 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 incident:

- Identifies victim workloads whose response time has crossed the performance threshold, which is the upper boundary of the expected range.
 - For volumes on HDD or Flash Pool (hybrid) aggregates, incidents are triggered only when the response time is greater than 5 milliseconds (ms) and the operations are more than 10 operations per second (ops/sec).
 - For volumes on all-SSD aggregates, incidents are triggered only when the response time is greater than 1 ms and the operations are more than 100 ops/sec.
- Identifies the cluster component in contention.

Note: If the response time of victim workloads at the cluster interconnect is greater than 1 ms, Performance Manager treats this as significant and triggers an incident 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 incident might occur for only a brief moment and then correct itself after the component it is using is no longer in contention. A continuous incident is one that reoccurs for the same cluster component within a five-minute interval and remains in the new state. For continuous incidents, Performance Manager triggers an alert after detecting the same incident during two consecutive analysis intervals. Incidents that remain unresolved, which have a state of new, can display different description messages as workloads involved in the incident change.

When an incident is resolved, it remains available in Performance Manager as part of the record of past performance issues for a volume. Each incident has a unique ID that identifies the incident type and the volumes, cluster, and cluster components involved.

Note: A single volume can be involved in more than one incident at the same time.

Incident state

Incidents can be in one of the following states:

New

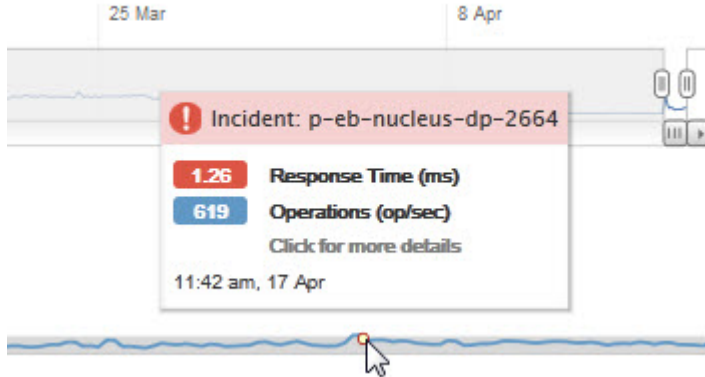
Indicates that the incident has not corrected itself, or has not been resolved, and the I/O response time of the impacted workloads remains above the performance threshold of the expected range.

Obsolete

Indicates that the incident has corrected itself, or has been resolved, and the I/O response time of the impacted workloads is no longer above the performance threshold of the expected range. A user might have made a change to the cluster to resolve the incident or the incident might have corrected itself by returning back within the expected range.

Incident notification

The incident alerts are displayed on the Dashboard, Volume Details page, and are sent to specified email addresses. If you have configured OnCommand Unified Manager to receive incident alerts from Performance Manager, the incidents are also displayed on the Unified Manager Dashboard. You can view detailed analysis information about an incident and get suggestions for resolving it on the Incident Details page.



In this example, an incident is indicated by a red dot (●) on the Response Time chart on the Volume Details page. Hovering your mouse cursor over the red dot displays a popup with more details about the incident and options for analyzing it.

Incident interaction

On the Volume Details page, you can interact with incidents in the following ways:

- Moving the pointer over a red dot displays a message that shows the incident ID, along with the response time, number of operations per second, and the date and time when the incident was detected.

If there are multiple incidents for the same time period, the message shows the number of incidents, along with the average response time and operations per second for the volume.

- Clicking a single incident displays a dialog box that shows more detailed information about the incident, including the cluster components that are involved, similar to the Summary section on the Incident Details page.

The component in contention is circled and highlighted red. You can click either the incident ID or **View full analysis** to view the full analysis on the Incident Details page. If there are multiple incidents for the same time period, the dialog box shows details about the three most recent incidents. You can click an incident ID to view the incident analysis on the Incident Details page. If there are more than three incidents for the same time period, clicking the red dot does not display the dialog box.

Related concepts

[Collecting data and monitoring workload performance](#) on page 46

[How Performance Manager determines the performance impact for an incident](#) on page 62

[Roles of workloads involved in a performance incident](#) on page 65

[Why a cluster component can be in contention](#) on page 63

[Types of workloads monitored by Performance Manager](#) on page 47

Related tasks

[Displaying information about a performance incident](#) on page 87

[Identifying victim workloads involved in a performance incident](#) on page 88

[Identifying bully workloads involved in a performance incident](#) on page 89

Related references

[Incident Details page](#) on page 108

How Performance Manager determines the performance impact for an incident

Performance Manager uses the deviation in activity, utilization, write throughput, cluster component usage, or I/O response time for a workload to determine the level of impact to workload performance. This information determines the role of each workload in the incident and how they are ranked on the Incident Details page.

Performance 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 incident.

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 response time is usually around 7 ms/op. The expected range for Workload B is 10-20 ms/op and its actual response time is usually around 15 ms/op. Both workloads are well within their expected range for response time. Due to contention on the cluster, the response time of both workloads increases to 40 ms/op, crossing the performance threshold, which is the upper bounds of the expected range, and triggering incidents. The deviation in response time, 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 response time of both workloads spike to 40 ms/op, but Workload A had the bigger performance impact because it had the higher response time deviation at 33 ms/op.

On the Incident Details page, in the Workload Details table, you can sort workloads by their deviation in activity, utilization, or throughput for a cluster component. You can also sort workloads by response time. When you select a sort option, Performance Manager analyzes the deviation in activity, utilization, throughput, or response time since the incident was detected from the expected values to determine the workload sort order. For the response time, the red dots (●) indicate a performance threshold crossing by a victim workload, and the subsequent impact to the response time. Each red dot indicates a higher level of deviation in response time, which helps you identify the victim workloads whose response time was impacted the most by an incident.

Related concepts

[Collecting data and monitoring workload performance](#) on page 46

[Roles of workloads involved in a performance incident](#) on page 65

[Why a cluster component can be in contention](#) on page 63

Types of workloads monitored by Performance Manager on page 47

Related tasks

Displaying information about a performance incident on page 87

Identifying victim workloads involved in a performance incident on page 88

Identifying bully workloads involved in a performance incident on page 89

Related references

Performance incident analysis and notification on page 59

Incident Details page on page 108

Why a cluster component 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 for client requests increases, which triggers an incident in Performance 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 response time. 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 Performance Manager collects and analyzes workload activity 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.

Performance 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 protocols (FCP) on the cluster. The wait time is time spent waiting for iSCSI Ready to Transfer (R2T) or FCP Transfer Ready (XFER_RDY) transactions to complete 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 response time 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 incident was detected. If the network processing component is in contention, it means high utilization at the network processing node is impacting the response time of one or more workloads.

Policy Group

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 response time 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 response time 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 incident was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the response time 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 response time 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

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 response time 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).

Note: When viewing SSD Aggregates, bully and shark workloads are not currently displayed, and utilization charts are unavailable.

Related concepts

[Collecting data and monitoring workload performance](#) on page 46

[Roles of workloads involved in a performance incident](#) on page 65

Types of workloads monitored by Performance Manager on page 47

Related tasks

Displaying information about a performance incident on page 87

Identifying victim workloads involved in a performance incident on page 88

Identifying bully workloads involved in a performance incident on page 89

Related references

Performance incident analysis and notification on page 59

Incident Details page on page 108

Roles of workloads involved in a performance incident

Performance Manager uses roles to identify the involvement of a workload in a performance incident. The roles include victims, bullies, and sharks. A user-defined workload can be a victim, bully, and shark at the same time.

The following table defines the workload roles:

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. Performance Manager identifies victim workloads based on their deviation in response time, where the actual response time, during an incident, has greatly increased from its expected range of response time.
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. Performance Manager identifies bully workloads based on their deviation in usage of a cluster component, where the actual usage, during an incident, 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 incident. Performance Manager identifies shark workloads based on their usage of a cluster component during an incident.

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 is unable to 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 snapshots, can also escalate into "bullying."

When Performance Manager detects an incident, it identifies all workloads and cluster components involved, including the bully workloads that caused the incident, the cluster component that is in contention, and the victim workloads whose performance has decreased due to the increased activity of bully workloads.

Note: If Performance Manager is unable to identify the bully workloads, it only alerts on the victim workloads and the cluster component involved.

Performance 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 incident might change its role or no longer be a participant in the incident. On the 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.

Related concepts

[Cluster concepts](#) on page 24

[Why a cluster component can be in contention](#) on page 63

[Types of workloads monitored by Performance Manager](#) on page 47

Related tasks

[Displaying information about a performance incident](#) on page 87

[Identifying victim workloads involved in a performance incident](#) on page 88

[Identifying bully workloads involved in a performance incident](#) on page 89

Related references

[Performance incident analysis and notification](#) on page 59

[Incident Details page](#) on page 108

Cluster configuration changes detected by Performance Manager

Performance Manager monitors your clusters for configuration changes to help you determine if a change might have caused or contributed to a performance incident. The Volume Details page displays a change event icon to indicate the date and time when the change was detected.

You can review the performance charts on the Volume Details page to see if the change event impacted the performance of the selected volume workload. If the change was detected at or around the same time as a performance incident, the change might have contributed to the slowdown in response time, which caused the incident alert to trigger.

Performance Manager can detect the following change events:

- A volume moves between aggregates. Performance Manager can detect when the move is in progress, completed, or failed. If Performance Manager is down during a volume move, when

Performance Manager is back up it detects the volume move and displays a change event for it on the Volume Details page.

- The throughput 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 response time, which might also trigger incidents for the policy group. The response time gradually returns back to normal and any incidents caused by the spikes become obsolete.
- A node in an HA pair takes over or gives back the storage of its partner node. Performance Manager can detect when the takeover or giveback operation has completed. If the takeover is caused by a panicked node, Performance Manager does not detect the event.
- A Data ONTAP upgrade or revert operation successfully completes. The previous version and new version are displayed.

Related concepts

[*How moving a FlexVol volume works*](#) on page 81

[*How the maximum throughput limit works*](#) on page 35

[*What an HA pair is*](#) on page 25

Navigating Performance Manager

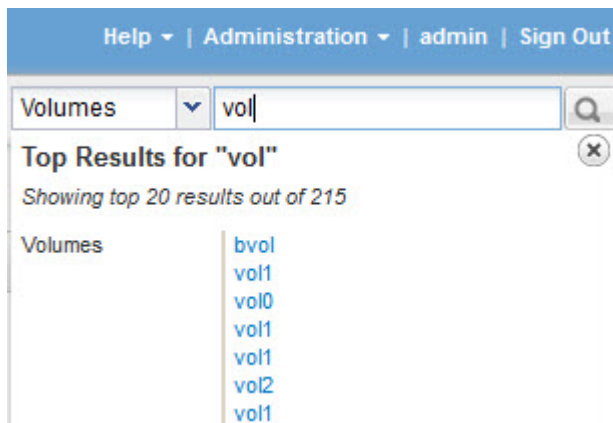
Performance Manager contains three main pages: Dashboard, Volume Details page, and Incident Details page. You use tabs and links to navigate Performance Manager. The tabs are across the top of the interface.

Note: A page in Performance Manager might display a large amount of information. To see all of the available information, always scroll to the bottom of the page.

Object search

To quickly access a specific object you want to analyze, you can use the Search bar at the top-right of the interface. You select the object type, including volume, LUN, or incident, and then enter three or more characters of the object name to display the search results, such as:

- volumes - volume name
- LUNs - LUN path
- incidents - incident ID



In the example above, the Search bar has the Volumes object type selected. Typing "vol" displays a list of all volume workloads whose names contain these characters. Below the list of volumes is a list of the associated LUN names, if available. Searching for a LUN displays a list of the associated volume names. If you upgrade Performance Manager, you can no longer search for incidents that occurred prior to the upgrade.

Note: If the search results display several volumes with the same name, the name of the associated clusters and SVMs are not displayed.

Linked object names

On the Incident Details page, you can click a workload name link to display the workload on the Volume Details page. On the Volume Details page, you can click an incident ID link to display the incident on the Incident Details page.

Incident: p-eb-nucleus-ag-31905 ⓘ
On Cluster: nucleus

Summary

Detected: 12:01 am, 13 Jun
State: Obsolete - 12:11 am, 13 Jun ⓘ
Description: [voX](#) is slow at aggr2

Component in Contention:

Network Network Processing Policy Group Cluster Interconnect Data Processing **Aggregate (aggr2)**

In the example above, on the Incident Details page, an incident involves a single volume workload. The volume name is a link you can click to display the volume on the Volume Details page. If an incident involves more than one volume, the number of volumes is displayed, but the volumes do not display a link.

Related tasks

[Searching for storage objects](#) on page 121

Related references

[Dashboard details for Incidents](#) on page 107

Logging in to the GUI

You can log in to the GUI using a supported web browser. You will be automatically logged out of the session after 4 hours.

Before you begin

Ensure that the web browser meets the minimum requirements.

Steps

1. Enter the following URL in your web browser: `http://URL`, where *URL* is the IP address or fully-qualified domain name (FQDN) of the virtual machine (VM) where the OnCommand application is running.
2. At the login screen, enter your user name and password.

Related concepts

[Browser and platform requirements](#) on page 70

Browser and platform requirements

To use the Performance Manager GUI, you must use a supported browser that runs on a supported client platform.

Performance Manager has been tested with the following browsers and client platforms; other browsers might work but have not been qualified. See the Interoperability Matrix at mysupport.netapp.com/matrix for the complete list of supported browser versions.

Supported browsers

- Mozilla Firefox versions 24 and 31
- Microsoft Internet Explorer (IE) versions 10 and 11
- Google Chrome version 36
- Apple Safari version 7

For IE, ensure that Compatibility View is disabled and Document Mode is set to the default. See the Microsoft IE documentation for information on these settings. For all browsers, disabling any popup blockers allows software features to display properly.

Supported browser client platforms

- Windows 7 and Windows XP
- Macintosh OS X 10.8 and 10.9

How graphs of performance data work

Performance 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 metrics 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 Volume Details page, the Historic data chart enables you to select a time frame for all graphs on the page. The 1d, 5d, 10d, 30d, 45d, and 90d buttons specify 1 day through 90 days (3 months) and the **Custom** button enables you to specify a custom time range within that 90 days. Each point on a graph represents a 5-minute collection interval, and a maximum of 90 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

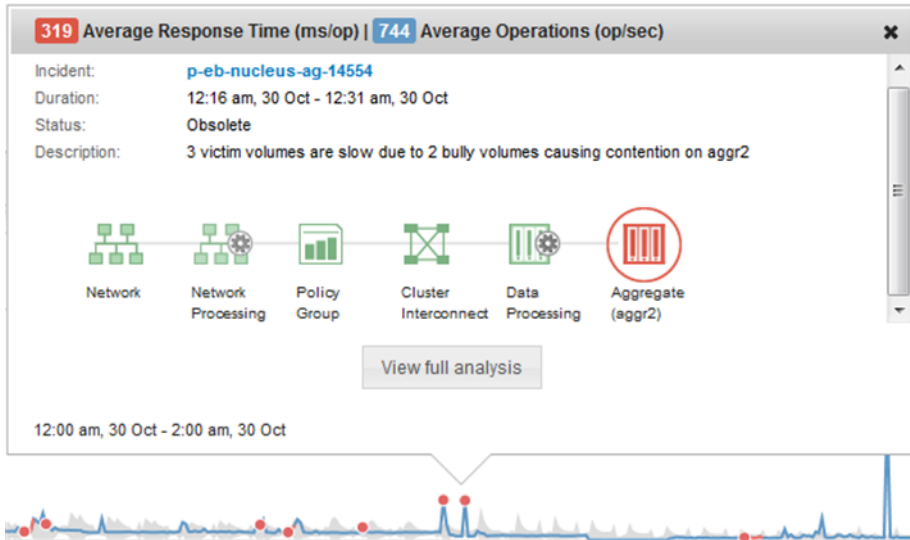
To view data point information on a graph, you can position the mouse cursor over a specific point within the graph, and a pop-up box displays listing the value and date and time information.



In this example, positioning the mouse cursor over the Operations chart on the Volume Details page displays the response time and operations values between 3:50 a.m. and 3:55 a.m. on October 20th.

Viewing event information

To view event information on a graph, you can position your mouse pointer 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 Volume Details page, clicking an incident event icon on the Response Time chart displays detailed information about the event in a pop-up box. The event is also highlighted in the Events List.

Related concepts

[Navigating Performance Manager](#) on page 68

Related references

[Incident Details page](#) on page 108

[Volume Details page](#) on page 82

[Performance statistics displayed in the data breakdown charts](#) on page 84

Exporting data to a CSV file

You can export information about all data displayed on the page you are viewing to a comma-separated variable (CSV) file. You can then open the file in a text editor or spreadsheet application.

Before you begin

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

About this task

Data points for every section currently displayed on the page are included in the CSV file. For example, on the Volume Details page, all data points in the Response Time chart and Operations chart will be included in the file, along with the data points for any Data breakdown charts.

Steps

1. Select **Actions > Export to CSV**.
2. In the dialog box, select to save the CSV file to a local or network location or open it in a compatible application.
3. Click **OK**.

Copying a link to a page

You can copy the link (URL) to the current page in your browser so that you can paste it into an email or another application and share it with other people. If you view the page for this URL at a later time, the page might display different data.

Before you begin

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

About this task

The link contains parameters so that all sections currently displayed on the page are included. For example, if any Data breakdown charts are currently displayed, the link will include parameters for each chart. When someone clicks the link, the Volume Details page is displayed in their browser just as it appeared when you copied the link.

Steps

1. Select **Actions > Copy link to page**.

The link for the page is displayed in a message window and is highlighted.

2. In the message window, press Control+C to copy the highlighted link.

The link is copied to your clipboard. You can paste the link in an email or a new browser window.

Printing a page

You can send a GUI page to a printer. The print action automatically formats the GUI page to display correctly in the print output.

Before you begin

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

Steps

1. Select **Actions > Print page**.
2. In the dialog box, select a printer and other settings.
3. Click **OK**.

Analyzing workload performance

Performance 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.

Related concepts

[Collecting data and monitoring workload performance](#) on page 46

[Types of workloads monitored by Performance Manager](#) on page 47

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

Related references

[Volume Details page](#) on page 82

Determining whether a workload has a performance issue

You can use OnCommand Performance Manager to determine whether a detected performance incident was truly caused by a performance issue on the cluster. The incident might have been caused a spike in activity, for example, that drove up its response time, but now the response time has returned to its usual levels.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You have identified the name of the volume, or associated LUN, you want to analyze.
- Performance Manager has collected and analyzed a minimum of five days of performance statistics from the cluster.

About this task

If you are viewing the Incident Details page, you can click the name link for a volume to go directly to the 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 Volume Details page.

3. In the **Historic data** chart, click **5d** to display the last five days of historical data.
4. Review the **Response Time** chart to answer the following questions:
 - Are there new performance incidents?
 - Are there historic performance incidents, 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?
 - Have there been configuration changes on the cluster that might have impacted performance?

If the response time for the volume does not display performance incidents, 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.

Related concepts

[Types of workloads monitored by Performance Manager](#) on page 47

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

[How graphs of performance data work](#) on page 70

Related tasks

[Searching for storage objects](#) on page 121

Related references

[Volume Details page](#) on page 82

Investigating a perceived slow response time for a workload

You can use OnCommand Performance Manager to determine if operations on the cluster might have contributed to the slow response time for a volume workload.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- You have identified the name of the volume, or associated LUN, you want to analyze.
- Performance Manager has collected and analyzed a minimum of five days of performance statistics from the cluster.

About this task

If you are viewing the Incident Details page, you can click the name for a volume to go directly to the 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 Volume Details page.

3. On the Historic data chart, click **5d** to display the last five days of historical data.

4. Review the **Operations** 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 **Operations**, select **Reads/writes/other**.

6. Click **Submit**.

The Reads/writes/other chart is displayed below the Operations 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.

Related concepts

[Types of workloads monitored by Performance Manager](#) on page 47

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

[How graphs of performance data work](#) on page 70

Related tasks

[Determining whether a workload has a performance issue](#) on page 75

[Searching for storage objects](#) on page 121

Related references

[Volume Details page](#) on page 82

Identifying trends of I/O response time on cluster components

You can use OnCommand Performance 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 have identified the name of the volume or associated LUN you want to analyze.
- To display 30 days of performance statistics, Performance Manager has 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 if the cluster is being overused or underused.

If you are viewing the Incident Details page, you can click the name for a volume to go directly to the 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 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 **Response Time**, select **Cluster Components** and **Reads/writes latency**.
6. Click **Submit**.

Both charts are displayed below the Response Time chart.

7. Review the **Cluster Components** chart. The chart breaks down the total response time by cluster component.

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 **Response Time** chart.

The Response Time 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 **Response Time** 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.

Related concepts

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

[Roles of workloads involved in a performance incident](#) on page 65

Related references

[Performance statistics displayed in the data breakdown charts](#) on page 84

Analyzing the performance improvements achieved from moving a volume

You can use OnCommand Performance Manager to investigate the impact of a volume move operation on the 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 have identified the name of the volume, or associated LUN, you want to analyze.
- Performance Manager has collected and analyzed seven days of data.

About this task

Performance Manager identifies when a volume moves between aggregates. It can detect when the volume move is occurring, completed, or failed. The 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 incident.

If you are viewing the Incident Details page, you can click the name of a volume to go directly to the 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 Volume Details page.


3. In the **Historic data** chart, adjust the sliders to show activity from the previous work week.

4. Analyze the **Response Time** chart and the **Operations** 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 incidents, 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 completed, you can return to Performance Manager to confirm that the move operation completed successfully and that the response time has improved.

5. If the **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 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 completed. A black, vertical line is also displayed from the change event icon to the Response Time 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 **Throughput**, 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 response time decreased significantly from over 42 ms/op to around 24 ms/op. The current response time 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.

Related concepts

[*Collecting data and monitoring workload performance*](#) on page 46

[*What performance events are*](#) on page 58

[*Cluster configuration changes detected by Performance Manager*](#) on page 66

[*Analyzing performance incidents*](#) on page 87

Related references

[*Performance statistics displayed in the data breakdown charts*](#) on page 84

[*Performance incident analysis and notification*](#) on page 59

How moving a FlexVol volume works

FlexVol volumes are moved from one aggregate or node to another within the same Storage Virtual Machine (SVM) for capacity utilization and improved performance, and to satisfy service-level agreements.

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 45 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.

You can change the amount of time client access is blocked or the number of times (*cutover attempts*) the final phase of the volume move operation is run if the defaults are not adequate. You can also determine what the system does if the volume move operation cannot be completed during the time client access is blocked. The `volume move start` man page contains details about moving a volume without disrupting client access.

Volume Details page

This page provides detailed performance statistics for all I/O activity and operations for the selected FlexVol volume workload. You can select a specific time frame over which to view the statistics and events for the volume. The events identify performance incidents and changes that might be impacting I/O performance.

Last Updated displays the date and time when you last refreshed the page in your browser. 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 operations across the number of days.

Note: When accessing the Performance Manager GUI from a Unified Manager GUI or email alert, you cannot access the Performance Manager Dashboard. Also, the Administration menu and Change Password option are not displayed.

Volume information

Information about the selected volume.

Volume

Displays the name and state of the volume. The volume state is displayed in parentheses. The state can be Online, Offline, Restricted, or Mixed. The Mixed state indicates that the constituents of an Infinite Volume are not all in the same state. See [Volume states](#) on page 33.

You can point the cursor to the volume name to display the full volume name and the names of the Storage Virtual Machine (SVM) and cluster that contain the volume.

Aggregate

Displays the name of the aggregate that the selected volume is on. You can point your cursor to the name to display details about the aggregate, including its state, capacity state, and capacity graphs for the data and Snapshot copies. See [Aggregate states](#) on page 43 and [Aggregate capacity states](#) on page 44.

Data

Displays the total data capacity and the used data capacity of the aggregate. If the aggregate is thin provisioned, which is also referred to as *aggregate overcommitment*, a flag is displayed with the overcommitted capacity. For more information about thin

provisioning, see [Considerations for using thin provisioning with FlexVol volumes](#) on page 34.

Snapshot Copies

This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

In both graphs, the capacity includes the Snapshot capacity when the used Snapshot capacity exceeds the Snapshot reserve.

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. Performance Manager retains a maximum of 90 days of historical data.

Note: 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 operations 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 90 days (**90d**), 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.

Note: If you access the Volume Details page by clicking the name link of a volume on the Incident 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 Volume Details page to display more detailed performance statistics for the selected volume.

Related concepts

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

[How cluster operations can affect workload response times](#) on page 54

[What the expected range of performance is](#) on page 50

What performance events are on page 58

Related tasks

Determining whether a workload has a performance issue on page 75

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 response time, 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 response time and operations across the number of days.

Note: You can use the Select All checkbox to select, or de-select, all the listed chart options.

Response Time

The following charts detail the response time data 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 response time impact by each component as it relates to the total response time. You can use the check box next to each component to show and hide its graph.

For 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 response times 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 response time for the selected volume workload, not all workloads on the cluster.

Note: The read and write statistics might not always add up to the total response time statistics displayed in the Response Time chart. This is expected behavior based on how Performance Manager collects and analyzes read and write statistics for a workload.

Policy Group Impact

Displays a graph of the percentage of the response time for the selected volume workload that is impacted by the throughput limit on its policy group.

If the workload is throttled, the percentage indicates how much the throttling contributed to the response time 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.

Note: The chart displays for workloads in a user-defined 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 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 Performance Manager, the last modified time is the date and time when Performance Manager first discovered the cluster.

Operations

The following charts detail the operations data for the selected workload:

Reads/writes/other

Displays a graph showing the number of read and write operations and other operations, per second, over the selected time frame.

Other operations 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 operations 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.

Throughput

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 Pools. 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 OnCommand Unified Manager and OnCommand System Manager. For information about how a volume uses cache storage, see the *Clustered Data ONTAP Logical Storage Management Guide*.

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 if 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 Response Time chart, or the Reads/writes/other chart under the Operations 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.

Related concepts

[Cluster concepts](#) on page 24

Analyzing performance incidents

You can analyze performance incidents to identify when they were detected, whether they are still new, the workloads and cluster components involved, and the options for resolving the incidents on your own.

Related concepts

[*Why a cluster component can be in contention*](#) on page 63

[*How Performance Manager determines the performance impact for an incident*](#) on page 62

Related tasks

[*Displaying information about a performance incident*](#) on page 87

[*Identifying victim workloads involved in a performance incident*](#) on page 88

[*Identifying bully workloads involved in a performance incident*](#) on page 89

[*Identifying shark workloads involved in a performance incident*](#) on page 91

Related references

[*Performance incident analysis and notification*](#) on page 59

[*Incident Details page*](#) on page 108

[*Volume Details page*](#) on page 82

Displaying information about a performance incident

You can use the Incident Details page in to analyze a specific incident and to view suggested actions for resolving it. You access the Incident Details page from email alerts, the Dashboard, or the Volume Details page.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There are new or obsolete performance incidents.

About this task

If you have connected OnCommand Performance Manager to OnCommand Unified Manager, performance incidents are displayed on the Dashboard in OnCommand Unified Manager. For information on configuring this connection, refer to the *OnCommand Performance Manager Installation and Administration Guide*.

Steps

1. On the **Dashboard**, locate a new incident you want to analyze. You can also type the name of an incident in the **Search** bar or click the incident name link in an incident email alert.

Under Filters, you can select specific clusters or detection times for which to display incidents.

2. Click the incident name.

The Incident Details page is displayed.

3. Record the incident ID number, located at the top of the **Incident Details** page, using an application such as a text editor.

You can search on this ID to locate the incident later.

Related concepts

[Analyzing performance incidents](#) on page 87

[Why a cluster component can be in contention](#) on page 63

Related references

[Performance incident analysis and notification](#) on page 59

[Dashboard details for Incidents](#) on page 107

Identifying victim workloads involved in a performance incident

In Performance Manager, you can identify which volume workloads have the highest deviation in response time caused by a cluster 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 are new or obsolete performance incidents.

About this task

The Incident 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 incident. The values are based on the peaks that Performance Manager identified when it detected and last analyzed the incident.

Steps

1. Display the **Incident Details** page to view information about the incident.

In the Workload Details table, the workloads are sorted on **Victims - Peak Deviation in Response Time**.

Note: When the table is sorted by peak deviation in response time, only user-defined workloads, such as volumes, are displayed. Workloads with very low response time values are not displayed in the table.

2. Click the name of a volume to view its current and historical response time and incidents on the **Volume Details** page.

Related concepts

[How Performance Manager determines the performance impact for an incident](#) on page 62

[Why a cluster component can be in contention](#) on page 63

Related tasks

[Displaying information about a performance incident](#) on page 87

[Identifying bully workloads involved in a performance incident](#) on page 89

[Identifying shark workloads involved in a performance incident](#) on page 91

Related references

[Incident Details page](#) on page 108

[Performance incident analysis and notification](#) on page 59

Identifying bully workloads involved in a performance incident

In Performance 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.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There are new or obsolete performance incidents.

About this task

The Incident 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 incident. The values are based on the peaks that Performance Manager identified when it detected and last analyzed the incident.

Steps

1. Display the **Incident Details** page to view information about the incident.
2. In the **Workload Details** table, select **Bullies - Peak Deviation in Activity** or **Bullies - Peak Deviation in Utilization**.

Note: By default, the table is sorted on **Victims - Peak Deviation in Response Time**. When the Workload details table is sorted by peak deviation in usage, workloads with low deviation in usage are not displayed in the table.

The workloads with the highest deviation in usage are displayed at the top of the table.

3. Click the name of a volume workload to view detailed information about its current and historical performance activity on the **Volume Details** page.

Related concepts

[How Performance Manager determines the performance impact for an incident](#) on page 62

[Why a cluster component can be in contention](#) on page 63

Related tasks

[Displaying information about a performance incident](#) on page 87

[Identifying victim workloads involved in a performance incident](#) on page 88

[Identifying shark workloads involved in a performance incident](#) on page 91

Related references

[Incident Details page](#) on page 108

[Performance incident analysis and notification](#) on page 59

Identifying shark workloads involved in a performance incident

In Performance Manager, you can identify which workloads have the highest deviation in usage for a cluster 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 or obsolete performance incidents.

About this task

The Incident 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 incident. The values are based on the peaks that Performance Manager identified when it detected and last analyzed the incident.

Steps

1. Display the **Incident Details** page to view information about the incident.
2. In the **Workload Details** table, select **Sharks - Peak Utilization** or **Sharks - Peak Activity**.

Note: By default, the table is sorted on **Victims - Peak Deviation in Response Time**. When the Workload details table is sorted by peak deviation in usage, workloads with low deviation in usage are not displayed in the table.

The workload with the highest usage for the component in contention, depending on the component type, is displayed at the top of the table.

3. Click the name of a volume workload to view detailed information about its current and historical performance activity on the **Volume Details** page.

Performance incident analysis for a MetroCluster configuration

You can use OnCommand Performance Manager to analyze a performance incident for a MetroCluster configuration. You can identify the workloads involved in the incident and review the suggested actions for resolving it.

MetroCluster incidents might be due to *bully* workloads that are over-utilizing the interswitch links (ISLs) between the clusters, or due to link health issues. Performance Manager monitors each cluster

in a MetroCluster configuration independently, without consideration of performance incidents on a partner cluster.

If you have connected Performance Manager to OnCommand Unified Manager, performance incidents from both clusters in the MetroCluster configuration are displayed on the Unified Manager Dashboard. For example, if you monitor each cluster using a separate instance of Performance Manager, you can connect each instance to Unified Manager to monitor performance incidents from one GUI. You can also use Unified Manager to check the health of each cluster and to view their relationship.

Related concepts

[Preparing for the MetroCluster installation](#) on page 37

[Performance monitoring of MetroCluster configurations](#) on page 55

[Performance incident analysis for a MetroCluster configuration](#) on page 91

[Roles of workloads involved in a performance incident](#) on page 65

Related references

[Performance incident analysis and notification](#) on page 59

Analyzing a performance incident on a cluster in a MetroCluster configuration

You can use OnCommand Performance Manager to analyze the cluster in a MetroCluster configuration on which a performance incident was detected. You can identify the cluster name, incident detection time, and the *bully* and *victim* workloads involved.

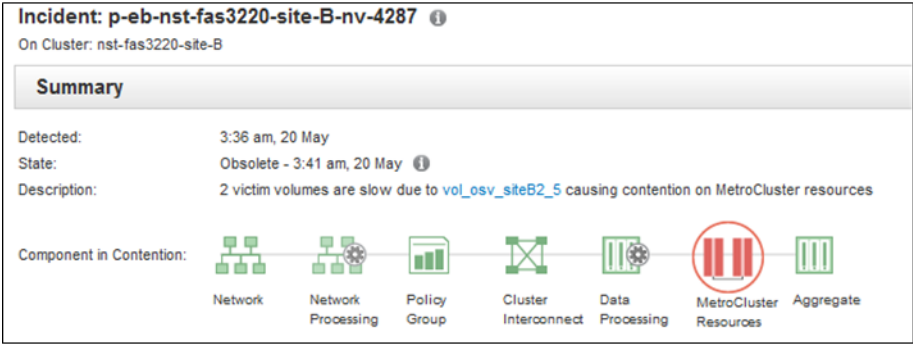
Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There are new or obsolete performance incidents for a MetroCluster configuration.
- Both clusters in the MetroCluster configuration are monitored by the same instance of Performance Manager.

Steps

1. Display the **Incident Details** page to view information about the incident.
2. Review the incident description to see the names of the workloads involved and the number of workloads involved:

Example

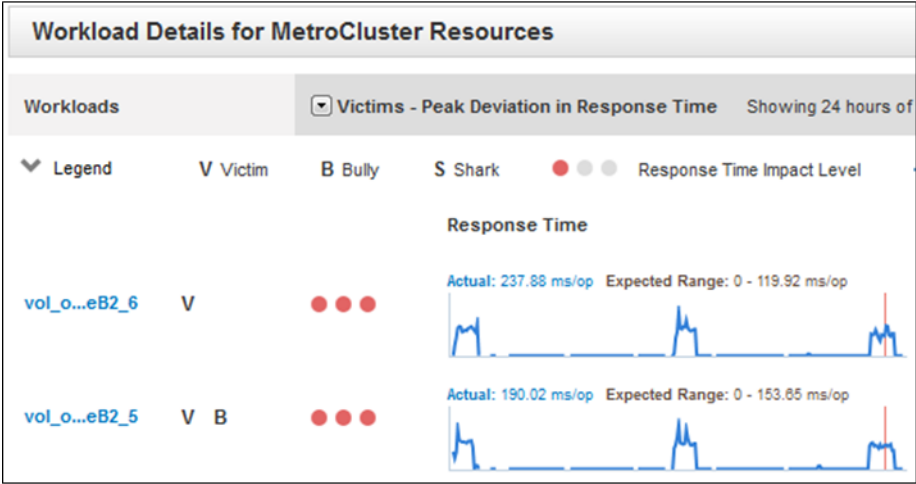


In this example, the MetroCluster Resources icon is red, indicating that the MetroCluster resources are in contention. You position your cursor over the icon to display a description of the icon. At the top of the page, **On Cluster** displays the name of the cluster on which the incident was detected.

- 3. Make a note of the cluster name and the incident detection time, which you can use to analyze performance incidents on the partner cluster.
- 4. In the **Workload Details** table, review the *victim* workloads to confirm that their response times are higher than the performance threshold.

Example

By default, the workloads are sorted by victims:



In this example, the victim workloads are displayed at the top of the table. The Response Time charts display, at a high-level, a consistent response time pattern for the victim workloads

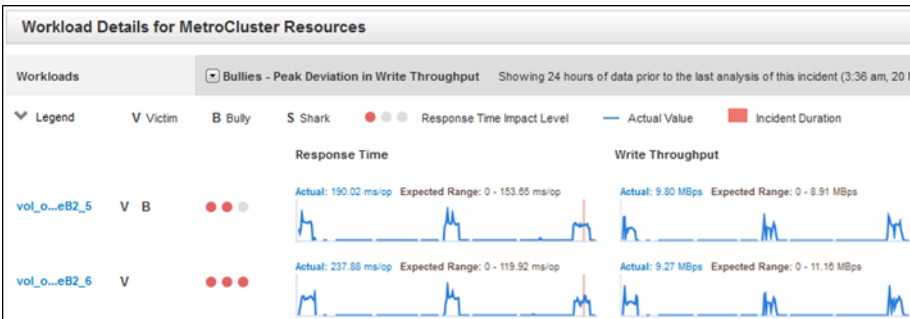
involved. Even though the abnormal response time of the victim workloads triggered the incident, a consistent response time pattern might indicate that the workloads are performing within their expected range, but that a spike in I/O increased the response time and triggered the incident.

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 response times to increase. If the response time for the workloads returns within the expected range, the incident state changes to obsolete, and remains in this state for more than 30 minutes, you can probably ignore the incident. If the incident is ongoing, and remains in the new state, you can investigate it further to determine if other issues caused the incident.

5. In the **Workload Details** table, select **Bullies - Peak Deviation in Write Throughput** to display the bully workloads at the top of the table.

Example

The bully workloads have a high deviation in write throughput:



In this example, one bully workload is displayed at the top of the table. The presence of bully workloads indicates that the incident might have been caused by one or more workloads on the local cluster overutilizing the MetroCluster resources.

The Total Write Throughput charts display, at a high-level, the write throughput pattern for the workloads. You can review the write throughput pattern to identify abnormal throughput, which might indicate that a workload is overutilizing the MetroCluster resources.

If no bully workloads are involved in the incident, the incident 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 OnCommand Unified Manager to check the health of both clusters in a MetroCluster configuration. You can also use Performance Manager to check for and analyze performance incidents on the partner cluster.

Related concepts

[Preparing for the MetroCluster installation](#) on page 37

[Performance incident analysis for a MetroCluster configuration](#) on page 91

Performance monitoring of MetroCluster configurations on page 55

Roles of workloads involved in a performance incident on page 65

Related tasks

Checking the health of clusters in a MetroCluster configuration on page 95

Analyzing a performance incident for a remote cluster on a MetroCluster configuration on page 98

Related references

Performance incident analysis and notification on page 59

Checking the health of clusters in a MetroCluster configuration

You can use OnCommand Unified Manager to check the operational health of clusters, and their components, in a MetroCluster configuration. If the clusters were involved in a performance incident, detected by OnCommand Performance Manager, the health status can help you determine if a hardware or software issue contributed to the incident.

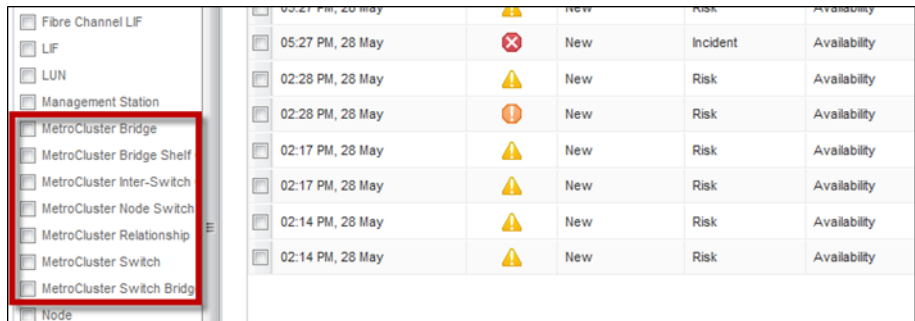
Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- In OnCommand Performance Manager, you have analyzed a performance incident for a MetroCluster configuration and obtained the name of the cluster involved.
- Both clusters in the MetroCluster configuration are monitored by the same instance of OnCommand Unified Manager.

Steps

1. Launch the OnCommand Unified Manager instance that is monitoring both clusters in the MetroCluster configuration involved in the performance incident.
2. Select **Events** to display the event list.
3. In the filter section, select all MetroCluster filters:

Example



<input type="checkbox"/>	Fibre Channel LIF				
<input type="checkbox"/>	LIF				
<input type="checkbox"/>	LUN				
<input type="checkbox"/>	Management Station				
<input type="checkbox"/>	MetroCluster Bridge				
<input type="checkbox"/>	MetroCluster Bridge Shelf				
<input type="checkbox"/>	MetroCluster Inter-Switch				
<input type="checkbox"/>	MetroCluster Node Switch				
<input type="checkbox"/>	MetroCluster Relationship				
<input type="checkbox"/>	MetroCluster Switch				
<input type="checkbox"/>	MetroCluster Switch Bridge				
<input type="checkbox"/>	Node				

<input type="checkbox"/>	05:27 PM, 28 May		New	Incident	Availability
<input type="checkbox"/>	02:28 PM, 28 May		New	Risk	Availability
<input type="checkbox"/>	02:28 PM, 28 May		New	Risk	Availability
<input type="checkbox"/>	02:17 PM, 28 May		New	Risk	Availability
<input type="checkbox"/>	02:17 PM, 28 May		New	Risk	Availability
<input type="checkbox"/>	02:14 PM, 28 May		New	Risk	Availability
<input type="checkbox"/>	02:14 PM, 28 May		New	Risk	Availability

In this example, the MetroCluster filters are outlined. You can select the filters so that only MetroCluster events are displayed in the event list.

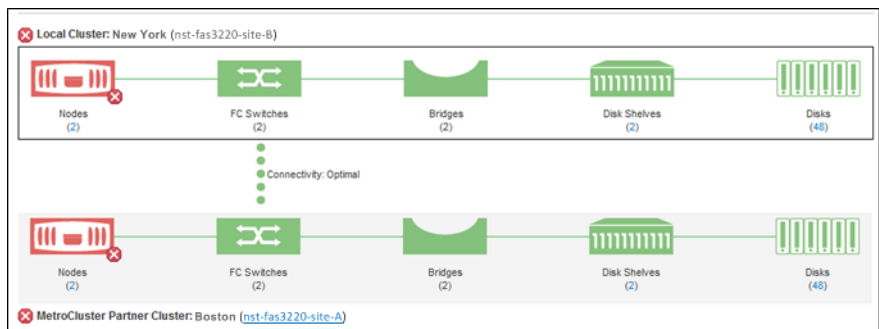
- Next to a MetroCluster event, click the name of the cluster.

The Clusters page is displayed with detailed information about the event.

Note: If no MetroCluster events are displayed, you can use the Search bar to search for the name of the cluster involved in the performance incident.

- Select the **MetroCluster Connectivity** tab to display the health of the connection between the selected cluster and its partner cluster.

Example

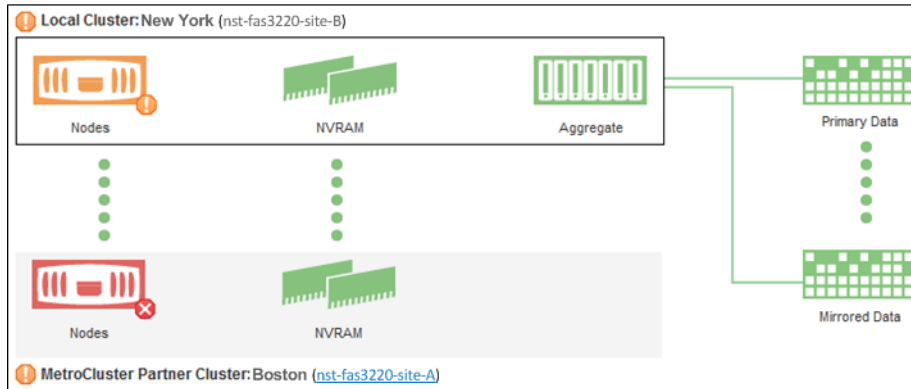


In this example, the names and the components of the local cluster and its partner cluster are displayed. A yellow or red icon indicates a health event for the highlighted component. The Connectivity icon represents the link between the clusters. You can point your mouse cursor to an icon to display event information or click the icon to display the events. A health issue on either cluster might have contributed to the performance incident detected in Performance Manager.

Performance 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 incident.

6. Select the **MetroCluster Replication** tab.

Example



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 incident. If there are no red or yellow icons on the page, a performance issue on the partner cluster might have caused the performance incident.

Related concepts

[Preparing for the MetroCluster installation](#) on page 37

[Performance monitoring of MetroCluster configurations](#) on page 55

[Roles of workloads involved in a performance incident](#) on page 65

Related tasks

[Analyzing a performance incident on a cluster in a MetroCluster configuration](#) on page 92

[Analyzing a performance incident for a remote cluster on a MetroCluster configuration](#) on page 98

Related references

[Performance incident analysis and notification](#) on page 59

Analyzing a performance incident for a remote cluster on a MetroCluster configuration

You can use OnCommand Performance Manager to analyze performance incidents on a remote cluster in a MetroCluster configuration. The analysis helps you determine if an incident on the remote cluster caused an incident on its partner cluster.

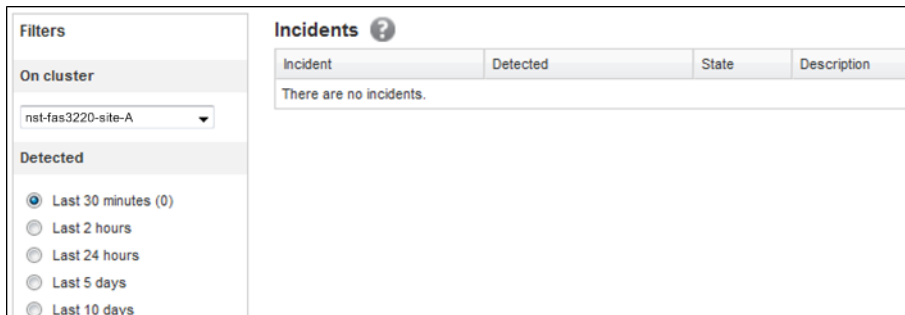
Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- In Performance Manager, you have analyzed a performance incident on a local cluster in a MetroCluster configuration and obtained the incident detection time.
- In OnCommand Unified Manager, you have checked the health of the local cluster and its partner cluster involved in the performance incident and obtained the name of the partner cluster.

Steps

1. Log in to the Performance Manager instance that is monitoring the partner cluster.
The Dashboard is displayed.
2. In the **Filters** section, for **On Cluster**, select the name of the partner cluster:

Example



The screenshot shows the OnCommand Performance Manager interface. On the left, the **Filters** section is expanded, showing the **On cluster** dropdown menu with 'nst-fas3220-site-A' selected. Below it, the **Detected** filter is set to 'Last 30 minutes (0)'. On the right, the **Incidents** table is displayed with columns: Incident, Detected, State, and Description. The table contains the message 'There are no incidents.'

In this example, the Incidents table lists the incidents for the cluster nst-fas3220-site-A. By default, Last 30 minutes is selected for the Detected filter. There no incidents for the selected cluster over the last 30 minutes, indicating that the cluster hasn't experience any performance issues during the time that the incident was detected on its partner.

3. If the selected cluster has incidents detected over the last 30 minutes, compare the incident detection time to the incident detection time for the incident on the local cluster.

If these incidents involve bully workloads causing contention on the data processing component, one or more of these bullies might have caused the incident on the local cluster. You can click the incident to analyze it and review the suggested actions for resolving it on the Incident Details page.

If these incidents do not involve bully workloads, they did not cause the performance incident on the local cluster.

Related concepts

[Preparing for the MetroCluster installation](#) on page 37

[Performance monitoring of MetroCluster configurations](#) on page 55

[Roles of workloads involved in a performance incident](#) on page 65

Related tasks

[Analyzing a performance incident on a cluster in a MetroCluster configuration](#) on page 92

[Checking the health of clusters in a MetroCluster configuration](#) on page 95

Related references

[Performance incident analysis and notification](#) on page 59

Responding to a performance incident caused by QoS policy group throttling

You can use OnCommand Performance Manager to investigate a performance incident caused by a policy group throttling workload throughput. The throttling increased the response times of volume workloads in the policy group. You can use the incident 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 are new or obsolete performance incidents.

Steps

1. Display the **Incident Details** page to view information about the incident.
2. Under **Summary**, read the **Description**, which displays the name of the workloads impacted by the throttling.

Note: The description can display the same workload for the victim and bully, because the throttling makes the workload a victim of itself. For a 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 Performance Manager, the last modified time is the date and time when Performance Manager first discovered the cluster.

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 at a later time.

4. In the **Workload Details** table, click **Bullies - Peak Deviation in Activity**.

The workloads in the policy group are sorted by highest deviation of actual activity from their expected activity. 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.

5. In the **Workloads** column, click the name of the top workload.

The Volume Details page is displayed, with detailed performance data for the selected workload.

6. Select **Break down data by**.

7. Select the check box next to **Response Time** to select all response time breakdown charts.

8. Under **Operations**, select **Reads/writes/other**.

9. Click **Submit**.

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

10. Compare the **Policy Group Impact** chart to the **Response Time** chart to see what percentage of throttling impacted the response time at the time of the incident.


The policy group has a maximum throughput of 1,000 operations per second (op/sec), which the workloads in it cannot collectively exceed. At the time of the incident, the workloads in the policy group had a combined throughput of over 1,200 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 response time, confirming that the throttling caused the incident to occur.

11. Review the **Cluster Components** chart, which shows the total response time by cluster component.

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

12. 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 response time. You can use these values when deciding to put a policy group limit on the throughput or operations.

13. Use OnCommand System Manager to increase the current limit on the policy group to 1,300 op/sec.
14. After a day, return to Performance Manager and search for the name of the workload that you recorded in Step 3.
The Volume Details page is displayed.
15. Select **Break down data by > Operations**.
16. Click **Submit**.
The Reads/writes/other chart is displayed.
17. At the bottom of the page, point your cursor to the change event icon () for the policy group limit change.
18. Compare the **Reads/writes/other** chart to the **Response Time** chart.
The read and write requests are the same, but the throttling has stopped and the response time has decreased.

Related concepts

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

[Why a cluster component can be in contention](#) on page 63

[Roles of workloads involved in a performance incident](#) on page 65

Related tasks

[Searching for storage objects](#) on page 121

Related references

[Performance statistics displayed in the data breakdown charts](#) on page 84

Responding to a performance incident caused by a disk failure

You can use OnCommand Performance Manager to investigate a performance incident caused by workloads overutilizing an aggregate. You can also use OnCommand Unified Manager to check the health of the aggregate to see if recent events detected on the aggregate contributed to the incident.

Before you begin

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

- There are new or obsolete performance incidents.

Steps

1. Display the **Incident Details** page to view information about the incident.
2. Under **Summary**, read the **Description**, which describes the workloads involved in the incident and the cluster component in contention.

There are multiple victim volumes, whose response time 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 Details** table, click **Bullies - Peak Deviation in Utilization** to sort the workloads on the aggregate by peak utilization.

The top workloads with the highest peak utilization since the incident was detected are displayed at the top of the table. One of the top workloads in the table 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 incident.

4. After confirming that the activity from the Disk Health workload caused the incident, wait for approximately 30 minutes for the reconstruction to finish and for Performance Manager to analyze the incident and detect whether the aggregate is still in contention.

5. In Performance Manager, search for the incident ID you recorded in Step 2.

The incident for the disk failure is displayed on the Incident Details page. After the RAID reconstruction is complete, under Summary, the Status is obsolete, indicating that the incident is resolved.

6. In the **Workload Details** table, click **Bullies - Peak Deviation in Utilization** to sort the workloads on the aggregate by peak utilization.

7. Click the name of a top volume workload.

Details for the selected volume are displayed on the Volume Details page.

8. Click **1d** to display the last 24 hours (1 day) of data for the selected volume.

In the Response Time chart, a red dot (●) indicates when the disk failure incident occurred.

9. Select **Break down data by**.

10. Under **Components**, select **Disk Utilization**.

11. 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.

12. Compare the data in the **Disk Utilization** chart to the data at the time of the incident in the **Response Time** chart.

At the time of the incident, the Disk Utilization shows a high amount of read and write activity, caused by the RAID reconstruction processes, which increased the response time of the selected volume. A few hours after the incident occurred, both the reads and writes and the response time have decreased, confirming that the aggregate is no longer in contention.

Related concepts

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

[Cluster configuration changes detected by Performance Manager](#) on page 66

[How Performance Manager determines the performance impact for an incident](#) on page 62

[Roles of workloads involved in a performance incident](#) on page 65

Related references

[Performance statistics displayed in the data breakdown charts](#) on page 84

Responding to a performance incident caused by HA takeover

You can use OnCommand Performance Manager to investigate a performance incident caused by high data processing on a cluster node that is in a high-availability (HA) pair. You can also use OnCommand Unified Manager to check the health of the nodes to see if any recent events detected on the nodes contributed to the incident.

Before you begin

- You must have the Operator, OnCommand Administrator, or Storage Administrator role.
- There are new or obsolete performance incidents.


Steps

1. Display the **Incident Details** page to view information about the incident.
2. Under **Summary**, read the **Description**, which describes the workloads involved in the incident and the cluster component in contention.

There is one victim volume, whose response time 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 incident is displayed in parentheses.

3. In the **Description**, click the name of the victim volume.

The Volume Details page is displayed. At the bottom of the page, in the Events time line, a change event icon () indicates the time that Performance Manager detected the start of the HA takeover.

4. Point your mouse cursor to the change event icon for the HA takeover.

Details about the HA takeover are displayed in the Events List table. In the Response Time chart, an incident indicates that the selected volume crossed the performance threshold due to high response time around the same time as the HA takeover.

5. Select **Break down data by**.

6. Under **Response Time**, select **Cluster Components**.

7. Click **Submit**.

The Cluster Components chart is displayed. The chart breaks down the total response time 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 response time for data processing to the total response time in the **Response Time** 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 response time and triggered the incident.

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 has completed, in Performance Manager, search for the incident ID you recorded in Step 2.

The incident triggered by the HA takeover is displayed on the Incident Details page. The incident now has a state of obsolete, which indicates that the incident is resolved.

12. In the **Description**, click the name of the victim volume.

The Volume Details page is displayed. At the bottom of the page, in the Events time line, a change event icon indicates the time that Performance Manager detected the completion of the HA giveback.

13. Select **Break down data by**.

14. Under **Response Time, 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 completed successfully.

16. In the **Cluster Components chart, compare the response time for data processing to the total response time in the **Response Time** chart.**

The response time at the data processing component has decreased, which has decreased the total response time. The node that the selected volume is now using for data processing has resolved the incident.

Related concepts

[How Performance Manager uses workload response time to identify performance issues](#) on page 53

[What an HA pair is](#) on page 25

[Cluster configuration changes detected by Performance Manager](#) on page 66

[How Performance Manager determines the performance impact for an incident](#) on page 62

[Roles of workloads involved in a performance incident](#) on page 65

Related references

[Performance statistics displayed in the data breakdown charts](#) on page 84

Page descriptions for analysis of performance incidents

You use the Dashboard to see a summary of incidents that recently occurred, or have been resolved, and the number of clusters and volumes impacted. You use the Incident Details page to see detailed information for a particular incident, which can be new or obsolete, and suggestions for resolving a new incident.

The topics below display when you click **Help** on the appropriate page.

Dashboard details for Quick Takes

The Quick Takes area enables you to view, as a graph, the clusters and volumes that have performance incidents and the number of incidents recently analyzed to be new or obsolete.

An incident can be in one the following states:

New

Indicates that the incident has not corrected itself, or has not been resolved, and the I/O response time of the impacted workloads remains above the performance threshold of the expected range.

Obsolete

Indicates that the incident has corrected itself, or has been resolved, and the I/O response time of the impacted workloads is no longer above the performance threshold of the expected range. A user might have made a change to the cluster to resolve the incident or the incident might have corrected itself by returning back within the expected range.

Clusters pane

Displays a bar chart that indicates the number of clusters with incidents whose state is new. The green bar displays the number of monitored clusters with no incidents. The red bar displays the number of clusters with new incidents. The total number of monitored clusters is displayed below the chart.

Note: If a cluster is unreachable, it is still included in the total number of monitored clusters.

Volumes pane

Displays a bar chart that indicates the volumes on all monitored clusters with incidents whose state is new. The green bar displays the number of volumes that do not have incidents. The red bar displays the number of volumes with new incidents. The total number of monitored volumes is displayed below the chart.

Note: If volumes are unreachable, they are still included in the total number of monitored volumes. Also, because Performance Manager does not monitor all cluster volumes (such as root volumes, volumes with LUNs in a policy group, or volumes associated with an SVM that is in a policy group), the number of volumes might not match the number of volumes on your clusters.

Recent pane

Displays the number of incidents whose state is new and the number of incidents over the last 24 hours whose state is obsolete.

Related concepts

[Why a cluster component can be in contention](#) on page 63

[Roles of workloads involved in a performance incident](#) on page 65

Related tasks

[Displaying information about a performance incident](#) on page 87

[Identifying victim workloads involved in a performance incident](#) on page 88

[Identifying bully workloads involved in a performance incident](#) on page 89

[Identifying shark workloads involved in a performance incident](#) on page 91

Related references

[Performance incident analysis and notification](#) on page 59

Dashboard details for Incidents

The Incidents table displays a list of performance incidents that Performance Manager detected to be new or obsolete within the last 30 minutes. The table provides you with a current list of incidents that might need your attention.

Incidents

Incident that recently became new are listed at the top of the table. By default, all incidents are sorted by their current status and detection time. You can use the filters to display incidents for all clusters, incidents for a specific cluster, or incidents that were detected over a specific period of time.

Incident

The incident ID. You can click the incident ID link to display the incident on the Incident Details page.

Detected

Time and date when Performance Manager detected the incident.

State

Current state of the incident, which can be one of the following values:

New

Indicates that the incident has not corrected itself, or has not been resolved, and the I/O response time of the impacted workloads remains above the performance threshold of the expected range.

Obsolete

Indicates that the incident has corrected itself, or has been resolved, and the I/O response time of the impacted workloads is no longer above the performance threshold of the expected range. A user might have made a change to the cluster to resolve the incident or the incident might have corrected itself by returning back within the expected range.

Description

Brief description of the incident, including the name or number of the workloads involved and the cluster component that is in contention.

Options**Filters**

- **On Cluster** provides a list of the clusters that Performance Manager is monitoring. You can select **All** to view incidents for all clusters. You can select the name of a cluster to display only the incidents for the cluster.

- **Detected** provides a time range during which incidents were detected. The list of incidents is updated for the selected time range.

Note: The selected filters might not update the incident values displayed in the Recent pane.

Related tasks

[Displaying information about a performance incident](#) on page 87

[Identifying victim workloads involved in a performance incident](#) on page 88

[Identifying bully workloads involved in a performance incident](#) on page 89

[Identifying shark workloads involved in a performance incident](#) on page 91

Related references

[Performance incident analysis and notification](#) on page 59

Incident Details page

This page displays detailed information about a new or obsolete performance incident for the selected cluster. It shows when the incident was detected, which workloads are involved, the cluster component in contention, and other information you can use to analyze and resolve the incident.

Last Updated displays the date and time when you last refreshed the page in your browser.

Note: When accessing the Performance Manager GUI from a Unified Manager GUI or email alert, you cannot access the Performance Manager Dashboard. Also, the Administration menu and Change Password option are not displayed.

Incident

Displays the ID of this incident and the name of the cluster on which the incident was detected; for example, p-eb-cluster01-np-1245. The following is the format of the ID, separated by hyphens:

- Incident type: p=performance
- Root cause type: eb=external bully volume, which is a user-defined workload
- Name of the cluster involved
- Cluster component in contention:
 - nw=network
 - pg=policy group
 - np=network processing
 - ci-cluster interconnect

- dp=data processing
- cpc=data processing incident due to bursts of write requests
- ag=aggregate made up of HDDs, or a mix of HDDs and SSDs
- ssd=aggregate made up of all SSDs (all-flash aggregate)
- nv=MetroCluster resources
- un=unknown
- Event ID: Numeric ID of the event stored in Performance Manager

You can make a note of the incident ID in case you want to search for it later.

Related concepts

Cluster concepts on page 24

Why a cluster component can be in contention on page 63

Types of workloads monitored by Performance Manager on page 47

Related tasks

Displaying information about a performance incident on page 87

Identifying victim workloads involved in a performance incident on page 88

Identifying bully workloads involved in a performance incident on page 89

Related references

Performance incident analysis and notification on page 59

What the incident Summary section displays

You use the Summary section on the Incident Details page to learn general information about a performance incident, including when it was detected, its current state, and a description of the workloads and cluster components involved.

Contents

Information about the workloads and the cluster components involved in the incident.

Detected

Date and time when Performance Manager detected the incident.

State

Current state of the incident. Possible values are:

New

Indicates that the incident has not corrected itself, or has not been resolved, and the I/O response time of the impacted workloads remains above the performance threshold of the expected range.

Obsolete

Indicates that the incident has corrected itself, or has been resolved, and the I/O response time of the impacted workloads is no longer above the performance threshold of the expected range. A user might have made a change to the cluster to resolve the incident or the incident might have corrected itself by returning back within the expected range.

Description

Brief description of the incident, including the name or number of each workload involved and the specific cluster component that is in contention. If a single workload is involved, you can click the name link to view details about it on the Volume Details page. You can point the cursor to the volume name to display the full volume name and the names of the Storage Virtual Machine (SVM) and cluster that contain the volume. For a 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 Performance Manager, the last modified time is the date and time when Performance Manager first discovered the cluster.

Note: Incidents that remain unresolved, which have a state of new, can display different description messages as workloads involved in the incident change.

Component in Contention

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 icons are displayed:

Network

Represents the wait time of I/O requests by the iSCSI protocols or the Fibre Channel protocols (FCP) on the cluster. The wait time is time spent waiting for iSCSI Ready to Transfer (R2T) or FCP Transfer Ready (XFER_RDY) transactions to complete 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 response time 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 incident was detected. If the network processing component is in contention, it means high utilization at the network processing node is impacting the response time of one or more workloads.

Policy Group

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 response time 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 response time 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 incident was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the response time 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 response time 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

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 response time 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).

Note: When viewing SSD Aggregates, bully and shark workloads are not currently displayed, and utilization charts are unavailable.

Related concepts

[Cluster concepts](#) on page 24

[Why a cluster component can be in contention](#) on page 63

[Types of workloads monitored by Performance Manager](#) on page 47

Related tasks

[Displaying information about a performance incident](#) on page 87

Related references

[Performance incident analysis and notification](#) on page 59

What the Workload Details table displays

On the Incident Details page, the Workload Details table lists a maximum of 50 workloads with the highest usage of the cluster component in contention or the highest response time. The performance statistics are based on the time the performance incident was detected up to the last time the incident was analyzed. The workloads might or might not be involved in the incident.

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 top-right of the table is the time and date range when an incident was detected and last analyzed. For a new incident, the last analyzed time will continue to update.

Note: If an incident was caused by a network component in contention, all affected workloads at that component are displayed in the table, and the table might take a long time to load. Also, you might not be able to export the page to a CSV file until the table has finished loading.

The top 5 workloads are displayed by default and you can click **Show next** to display additional workloads.

The table displays the following information:

Workloads column

Displays the name of the user-defined and system-defined workloads. For user-defined workloads, you can click the name link to display details about the incident on the Volume Details page. For system-defined workloads, such as storage efficiency and disk health processes, the type of workload is displayed. 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. Refer to [Workloads monitored by Performance Manager](#) on page 47.

An icon is displayed next to a workload name to indicate its role in the incident: "*sharks*," "*bullies*," or "*victims*."

You can point the cursor to the volume name to display the full volume name and the names of the Storage Virtual Machine (SVM) and cluster that contain the volume.

Workload Sort column

Enables you to sort the workloads based on their role in the incident, including "*sharks*," "*bullies*," or "*victims*," and displays details about their response time 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 response time or usage.

The data in the graphs shows 24 hours of performance statistics prior to the last time the incident was analyzed. The actual values and expected values for each workload are based on the time the workload was involved in the incident. For example, a workload might

become involved in an incident after the incident was detected, so its performance statistics might not match the values at the time of incident detection. By default, the workloads are sorted by peak (highest) deviation in response time, which displays the "victim" workloads at the top of the table. The column displays the following information:

Response Time column

Displays the impact of the incident to the response time of the workload during the last analysis. The actual response time is a blue line. A red bar highlights the incident duration, from the detection time to the last analyzed time.

The red dots indicate a performance threshold crossing, where the actual response time is above the upper bounds of the expected range. The number of red dots indicates the deviation of the actual response time from the expected response time since Performance Manager last analyzed the workload. The red dots and the victim icon identify "victim" workloads for the component that is involved in the incident. If there are red dots but no victim icon is displayed, either the workload response time increased due to contention on a different component or its response time and operations did not cross the static thresholds. If there are no red dots, the workload has not crossed the performance threshold.

For more information on the static thresholds, see [What incidents are](#) on page 59. For information about how Performance Manager ranks the workloads and determines the sort order, see [How Performance Manager determines the performance impact for an incident](#) on page 62.

Note: When you sort by peak deviation in response time, system-defined workloads are not displayed in the table, since response time applies only to user-defined workloads. Workloads with very low response time values are not displayed in the table.

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 incident duration, from the detection time to the last analyzed time. For more information, refer to [Workload performance measurements](#) on page 48.

Note: For the network component, because network performance statistics come from activity off the cluster, this column is not displayed.

Note: When the Workload details table is sorted by peak deviation in usage, workloads with low deviation in usage are not displayed in the table. 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.

For information about how Performance Manager ranks the workloads and determines the sort order, see [How Performance Manager determines the performance impact for an incident](#) on page 62.

Capacity column

Displays the total storage capacity, in gigabytes (GB), of the workload. The total capacity includes the entire footprint and allocated capacity. For example, if the allocated capacity is thin provisioned, the value includes the thin provisioned capacity. The footprint contains the user data and metadata, from processes such as deduplication and snapshots. For more information on volume footprints, see [What the volume footprint is](#) on page 33.

Related concepts

[Cluster concepts](#) on page 24

[How Performance Manager determines the performance impact for an incident](#) on page 62

[Why a cluster component can be in contention](#) on page 63

Related tasks

[Displaying information about a performance incident](#) on page 87

[Identifying victim workloads involved in a performance incident](#) on page 88

[Identifying bully workloads involved in a performance incident](#) on page 89

Related references

[Performance incident analysis and notification](#) on page 59

What Contention History charts display

On the Incident Details page, the Contention History charts display historical performance statistics for the cluster component in contention. You can review the charts to learn about the top victim workloads at the component, the history of utilization or activity for the component, and the history of incidents for the component.

The selected incident is indicated by a red-shaded line that appears in each of the displayed charts. The width of the red line indicates the incident duration. The incident duration is the time between detection, when the incident became new, and the time when the incident was last analyzed or, for an obsolete incident, when it was resolved.

For new incidents, the charts show 10 days of data prior to the incident detection. For obsolete incidents, the data prior to the incident detection is three times the incident duration. For example, if an incident was new for one day, and then became obsolete, the charts display three days of data before the incident detection and two days of data after the incident resolution. If there is not enough data before or after the incident, the charts display the maximum data available. If an incident has a very short duration of less than one hour, six hours of data is displayed before and after the incident. The data after the incident was resolved is two times the incident duration.

The following charts are displayed:

Top Victim Workloads

Displays the history of response time, in milliseconds per operation (ms/op), for the top victim workloads at the component in contention. The top victims have the highest peak deviation in response time. You can hide and show the statistics for a workload by

selecting the check box next to the workload name. You can point to the statistics for a displayed workload to view its response time at a specific point in time.

The response time values displayed in the Top Victim Workloads chart do not match the response time values displayed in the Cluster components breakdown chart on the Volume Details page. The Top Victim Workloads chart displays the actual response time at the cluster component, while the Cluster components breakdown chart provides a visualization of the amount of time spent at a cluster component. Also, since the Cluster components breakdown chart displays the combined response time values of components of the same type, these values do not match the values in the Top Victim Workloads chart, which shows the response time at the specific component in contention.

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 Throughput History

This chart displays for the MetroCluster Resources component only. It 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.

Incident History

Displays red-shaded lines to indicate the historic incidents for the component in contention. For obsolete incidents, the chart displays incidents that occurred before the selected incident was detected and after it was resolved.

Related concepts

[*Cluster concepts*](#) on page 24

[*How Performance Manager determines the performance impact for an incident*](#) on page 62

[*Why a cluster component can be in contention*](#) on page 63

[*Types of workloads monitored by Performance Manager*](#) on page 47

Related tasks

[*Displaying information about a performance incident*](#) on page 87

Related references

[*Performance incident analysis and notification*](#) on page 59

Suggested actions for resolving performance incidents

You can use the suggested actions to try and resolve performance incidents on your own. The first three suggestions are always displayed and the suggestions under the fourth suggestion are specific to the type of incident 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 involve using OnCommand Performance Manager, OnCommand System Manager, Data ONTAP commands, or a combination of these tools.

Related concepts

[How Performance Manager determines the performance impact for an incident](#) on page 62

Related references

[Performance incident analysis and notification](#) on page 59

Managing data sources

You can manage the Data ONTAP clusters you want to use in Performance Manager, including adding, editing, and removing clusters.

Adding clusters

You can add a cluster to Performance Manager to monitor the cluster and obtain information about its status and configuration.

Before you begin

- The clusters you want to add meet the configuration requirements and you have the required privilege for adding them.
- You must have the OnCommand Administrator role or the Storage Administrator role.

About this task

You cannot add the same cluster to more than one instance of Performance Manager. When you attempt to add a cluster that a Performance Manager instance is already monitoring, a warning message is displayed in the GUI. Each cluster in a MetroCluster configuration must be added separately.

A single instance of Performance Manager supports a specific number of clusters and volumes. If Performance Manager is monitoring an environment that exceeds the supported configuration, you might have difficulty collecting and analyzing configuration and performance data from the cluster. See the *OnCommand Performance Manager Release Notes* for the number of clusters, nodes, and volumes that Performance Manager can reliably support.

Steps

1. Click **Administration > Manage Data Sources**.
2. On the **Manage Data Sources** page, click **Add**.
3. In the **Add Cluster** dialog box, specify the values as required and then click **Save and Close**.

Result

The cluster is added to the Performance Manager database after the default monitoring interval of approximately 15 minutes. If you destroy a Performance Manager virtual machine (VM) and then install a new instance using the same IP address assigned to the previous VM, adding the clusters

from the previous VM to the new VM displays an error message that the clusters are already monitored. You can ignore this error message.

Note: If the UUID of a monitored cluster changes, due to a cluster rebuild, for example, Performance Manager does not associate the new UUID with the cluster and the cluster is no longer monitored. To associate the cluster to the new UUID, you must remove the cluster from Performance Manager and then re-add it.

Related tasks

[Controlling and monitoring I/O performance to FlexVol volumes by using Storage QoS](#) on page 36

Related references

[Performance Manager user roles and capabilities](#) on page 129

[Manage Data Sources page](#) on page 122

Requirements for adding a cluster to Performance Manager

You must have all necessary configuration information available before adding a cluster.

Required information

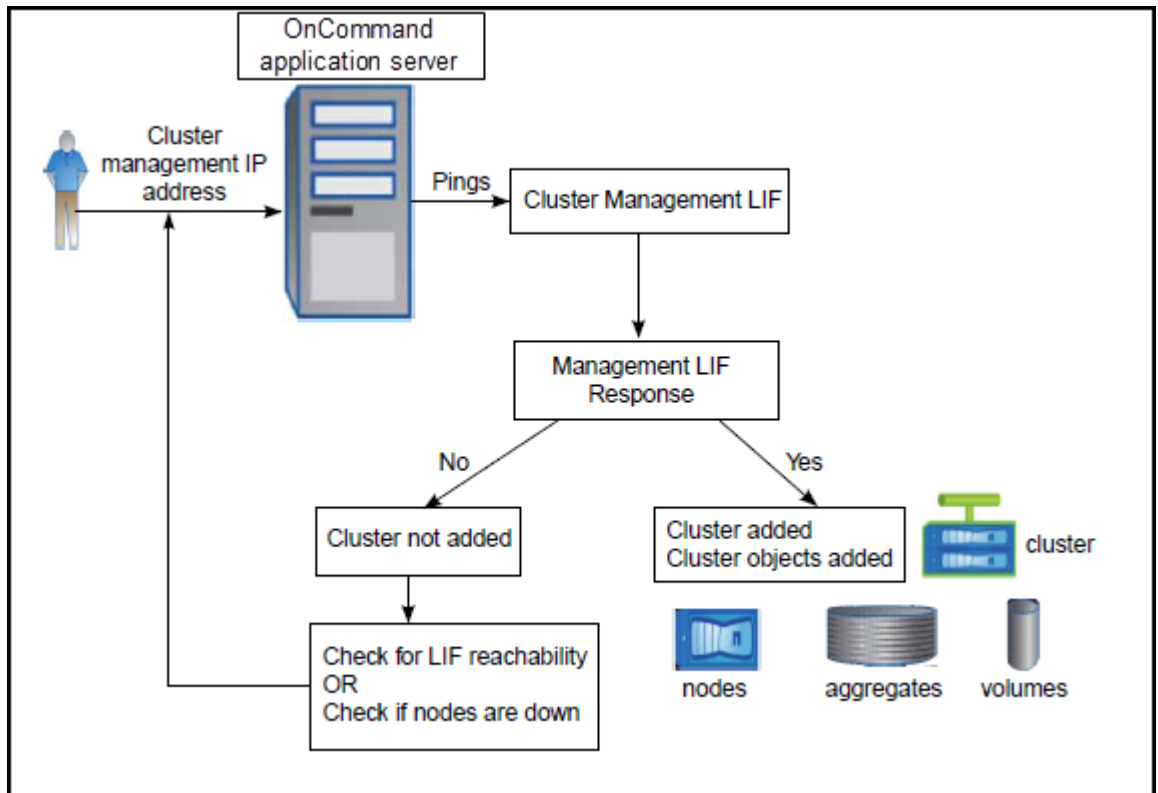
- Host name or cluster management IP address
The host name is the FQDN or short name that Performance Manager uses to connect to the cluster. This host name must resolve to the cluster management IP address.
The cluster management IP address must be the cluster management LIF. If you use a node management LIF, the operation fails.
- User name and password to access the cluster
If the version of Data ONTAP is earlier than 8.3, this account must have the *Admin* role with Application access set to *ontapi*.
- Type of protocol (HTTP or HTTPS) that is to be configured on the cluster and the port number of the cluster

How the discovery process works

After you have added the cluster to Performance 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 for cluster configuration information is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster details in the Performance Manager GUI.

The following image illustrates the discovery process:



Related tasks

[Adding clusters](#) on page 117

Viewing the clusters list

You can use the Manage 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.

Step

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

Related references

[Performance Manager user roles and capabilities](#) on page 129

[Manage Data Sources page](#) on page 122

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 Performance Manager. For example, you can change the protocol from HTTP to HTTPS using the Edit Cluster dialog box.

Before you begin

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

About this task

Attention: 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. Click **Administration > Manage Data Sources**.
2. On the **Manage 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 **Save**.

Related references

[Performance Manager user roles and capabilities](#) on page 129

[Manage Data Sources page](#) on page 122

Removing clusters

You can remove a cluster from Performance Manager by using the Manage Data Sources page. For example, you can remove a cluster 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

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

Steps

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

Result

The cluster, its storage objects along with the history, and all associated events are removed, and the cluster is no longer monitored by Performance Manager. The instance of Performance Manager registered with the removed cluster is also unregistered from the cluster.

Related references

[Performance Manager user roles and capabilities](#) on page 129

[Manage Data Sources page](#) on page 122

Searching for storage objects

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

Steps

1. Use the filter to select a specific storage object type for your search.
The available objects are All, Volumes, LUNs, or Incidents.
2. Type your search parameters into the search bar.
For example, type the name of the volume, or type any three characters in the volume's name.
If there are any matches to your search parameters, they are displayed in the Top Results drop-down list.
3. Select the object you want to view from the drop-down list.
The page for that object is displayed.

Related concepts

[Navigating Performance Manager](#) on page 68

Page descriptions for data source management

You can view and manage your clusters, including adding, editing, and removing clusters, from a single page.

The topics below display when you click **Help** on the appropriate page.

Manage Data Sources page

The Manage Data Sources page enables you to add clusters and to view detailed information about the clusters that you are monitoring.

Command buttons

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

Add

Opens the Add Cluster dialog box, which enables you to add clusters.

Edit

Opens the Edit Cluster dialog box, which enables you to edit the settings of the selected cluster.

Remove

Removes the selected cluster and all the associated events and storage objects. After the cluster is removed, it is no longer monitored.

Attention: The cluster, its storage objects, and all associated events are removed, and the cluster is no longer monitored by Performance Manager. The instance of Performance Manager registered with the removed clustered is also unregistered from the cluster.

Clusters list

The Clusters list displays, in tabular format, the properties of all the discovered clusters. You can use the column filters to customize the data that is displayed:

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 with port 443.

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.

Status

Displays the current status of the cluster:

- Normal—Cluster is operating normally.
- Authorization Failure—Invalid credentials for accessing the cluster.
- Network Access Failure—Network connection issue or a network timeout occurred.
- LIF Error—Issue with a node-management or cluster-management LIF.
- Collection Stopped—An older version and a newer version of Performance Manager are monitoring the same cluster, which has stopped data collection.
Only the newer version can monitor the cluster. You must remove the cluster from the older version for the newer version to resume data collection.
- Duplicate Datasource—Duplicate data sources, such as clusters, are being monitored.
- ZAPI Error—Issue with the cluster that is preventing data collection.
- Internal Error
- Unknown

Status Message

Brief description of the current cluster status.

Add Cluster dialog box

You can add an existing cluster to monitor the cluster and obtain information about its health, capacity, and configuration.

You can add a cluster by specifying the following options:

Host Name or IP Address

Enables you to specify the host name (preferred) or the IP address of the cluster-management 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 GUI, rather than trying to correlate an IP address on one page to a host name on another page, for example.

User Name

Enables you to specify a user name that can be used to log in to the cluster.

Password

Enables you to specify a password for the specified user name.

Protocol

Enables you to specify the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.

Port

Enables you to specify the port number of the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Related tasks

[Adding clusters](#) on page 117

Related references

[Performance Manager user roles and capabilities](#) on page 129

[Manage Data Sources page](#) on page 122

Edit Cluster dialog box

The Edit Cluster dialog box enables you to modify the settings of an existing cluster, including the IP address, port, and protocol. For example, you can change the protocol from HTTP to HTTPS.

You can edit the following fields:

Host Name or IP Address

Enables you to specify the FQDN, short name, or the IP address of the cluster-management LIF that is used to connect to the cluster.

User Name

Enables you to specify a user name that can be used to log in to the cluster.

Password

Enables you to specify a password for the specified user name.

Protocol

Enables you to specify the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.

Port

Enables you to specify the port number of the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Related tasks

[Editing clusters](#) on page 120

[Adding clusters](#) on page 117

Managing users

You can manage the users who use Performance Manager, including setting up user accounts, configuring user authentication, and assigning user roles for controlling access to specific features.

What the maintenance user does

Created during initial configuration, the maintenance user can create subsequent users and assign them roles. The maintenance user can also access the maintenance console and has the role of OnCommand Administrator in the GUI.

The maintenance user can perform the following functions using the maintenance console:

- Configure network access
- Upgrade to newer versions of Performance Manager
- Shut down virtual appliances (only from VMware console)
- Increase data disk or swap disk size
- Change the time zone
- Send on-demand AutoSupport messages to technical support from the maintenance console
- Generate support bundles to send to technical support

Related references

[*Definitions of user types*](#) on page 128

[*Definitions of user roles in Performance Manager*](#) on page 128

[*Performance Manager user roles and capabilities*](#) on page 129

What RBAC is

RBAC (role-based access control) provides the ability to control who has access to various features and resources in Performance Manager.

What RBAC 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.

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 Performance Manager.

The following LDAP servers are compatible with the management server:

- Microsoft Active Directory
- OpenLDAP
- IBM Lotus LDAP
- Netscape LDAP Server

You can use one of the following predefined authentication services or specify your own authentication service:

- Microsoft Active Directory

Note: 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 Performance 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 Performance 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.

Related tasks

[Enabling remote authentication](#) on page 138

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.

Your Performance Manager user types are as follows:

Maintenance user

Created from the maintenance console during the initial configuration of Performance Manager and its credentials are stored on the Performance Manager server. The maintenance user then creates additional users and assigns roles. The maintenance user is also the only user with access to the maintenance console.

Local user

Accesses the Performance Manager GUI using the credentials stored on the Performance Manager server. User perform functions based on the role given by the maintenance user or a user with the OnCommand Administrator role.

Remote group

Groups of users that access the Performance Manager GUI 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 Performance Manager GUI using their individual user credentials. Remote groups can perform functions according to their assigned roles.

Remote user

Accesses the Performance Manager GUI 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 Performance Manager database, has no access to the Performance Manager GUI or the maintenance console, and cannot execute API calls.

Related concepts

[What the maintenance user does](#) on page 126

Definitions of user roles in Performance Manager

The maintenance user or OnCommand administrator assigns a role to every user. Each role contains certain privileges. The scope of activities that you can perform in Performance Manager depends on the role you are assigned and which privileges the role contains.

The following predefined roles exist in Performance Manager:

Operator

Views storage system information and other data collected by Performance Manager.

Storage Administrator

Configures storage management operations within Performance Manager. The role enables the storage administrator to create alerts and configure other storage management-specific options.

OnCommand Administrator

Configures settings unrelated to storage management. The role enables the management of users, database access, and administrative options, including authentication, SMTP, networking, and AutoSupport.

Performance Manager user roles and capabilities

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

The following table displays the functions that each role can perform:

Function	Operator	Storage Administrator	OnCommand Administrator
View storage system information	•	•	•
View events	•	•	•
Define alerts		•	•
Manage storage management options		•	•
Manage users			•
Manage administrative options			•
Manage database access			•

Related concepts

[What the maintenance user does](#) on page 126

Adding users

You can create local users or database users from the Manage Users page. You can also add remote users or groups belonging to an authentication server. You can assign roles to these users, and based

on the privileges of the roles, users can effectively manage the storage objects and data using Performance Manager, or view data in a database.

Before you begin

- You must have the OnCommand Administrator role or the Storage Administrator role.
- To add a remote user or group, you must have enabled remote authentication and configured your authentication server.

About this task

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

Note: To ensure that you have at least one local user for accessing Performance Manager, you cannot delete the last local user administrator account.

Steps

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

Note: The specified email address must be unique to the instance of Performance Manager.

4. Click **Save and Close**.

Related tasks

[Enabling remote authentication](#) on page 138

[Setting up authentication services](#) on page 139

[Adding authentication servers](#) on page 141

Related references

[Performance Manager user roles and capabilities](#) on page 129

[Manage Users page](#) on page 133

Viewing users

You can use the Manage Users page to view the list of users who manage storage objects and data using Performance Manager. You can view details about the users, such as the name, type of user, email address, and role assigned to the users.

Step

1. Click **Administration > Manage Users**.

The list of users is displayed in the Manage Users page.

Related references

[Performance Manager user roles and capabilities](#) on page 129

[Manage Users page](#) on page 133

Editing the user settings

You can edit user settings, such as the email address and role specified for users on the Manage Users page. For example, you might want to change the role of a user who is a storage operator, and assign storage administrator privileges to that user.

Before you begin

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

About this task

When you modify the role assigned to a user, the changes are applied when either of the following occurs:

- The user logs out and logs back in to Performance Manager
- Session timeout of 24 hours has occurred

Steps

1. Click **Administration > Manage Users**.
2. In the **Manage Users** page, select the user that you want to edit and click **Edit**.
3. In the **Edit User** dialog box, edit the appropriate settings that are specified for the user.
4. Click **Save and Close**.

Related references

[Performance Manager user roles and capabilities](#) on page 129

[Manage Users page](#) on page 133

Changing the local user password

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

Before you begin

You are logged in as a local user.

About this task

The passwords for remote users or members of remote groups cannot be changed from the GUI. To change the remote user password, contact your password administrator.

Steps

1. Log in to the Performance Manager GUI.
2. Click ***user_name*** > **Change Password**.

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

3. In the **Change Password** dialog box, enter the details as required.
4. Click **Save**.

Related references

[Performance Manager user roles and capabilities](#) on page 129

Deleting users or groups

You can delete one or more users from the management server database to prevent the users from accessing Performance Manager.

Before you begin

- When you are deleting remote groups, you must first have reassigned the events that are assigned to users of those remote groups.
- You must have the OnCommand Administrator role or the Storage Administrator role.

About this task

Attention: To ensure that you have at least one user account for accessing Performance Manager, do not delete the maintenance user account.

Steps

1. Click **Administration > Manage Users**.
2. In the **Manage Users** page, select the users or groups that you want to delete and click **Delete**.
3. Click **Yes** to confirm the delete request.

Related references

[Performance Manager user roles and capabilities](#) on page 129

Page descriptions for user management

You can manage user access to Performance Manager, including adding, editing, and removing users, from a single page.

The topics below display when you click **Help** on the appropriate page.

Manage Users page

The Manage Users page displays a list of users and groups, and provides information such as the name, type of user, email address, and role. You can also 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, remote user, remote group, or a database user.

You can add remote users or groups only if your authentication server is enabled and configured.

Edit

Displays the Edit User dialog box, which enables you to edit the settings for the selected user.

Remove

Removes the selected users from the management server database.

List view

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

Name

Displays the name of the user or group.

Type

Displays the type of user. The user type can be Local User, Remote User, Remote Group, Database User, or Maintenance User.

Email

Displays the email address of the user.

Role

Displays the type of role that is assigned to the user. The role can be Operator, Storage Administrator, or OnCommand Administrator.

Note: This option is disabled for the Database User type.

Related tasks

[Viewing users](#) on page 131

[Adding users](#) on page 129

[Editing the user settings](#) on page 131

[Deleting users or groups](#) on page 132

Add User dialog box

You can create local users or database users, or add remote users or remote groups and assign roles so that these users can efficiently manage the storage objects and data using Performance 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 Performance Manager.

Password

Enables you to specify a password for the specified user name. This field is displayed only when you are adding a local user or a database user.

Confirm Password

Enables you to reenter your password to ensure the accuracy of what you entered in the Password field. This field is displayed only when you are adding a local user or a database user.

Email

Enables you to specify an email address for the user. This field is displayed only when you are adding a remote user or a local user.

Note: The specified email address must be unique to the instance of Performance Manager.

Role

Enables you to assign a role to the user and defines the scope of activities that the user can perform. The role can be OnCommand Administrator, Storage Administrator, or Operator.

Related tasks

[Adding users](#) on page 129

Related references

[Manage Users page](#) on page 133

[Performance Manager user roles and capabilities](#) on page 129

Edit User dialog box

The Edit User dialog box enables you to edit the email address or role of a selected user.

Details

The Details area enables you to edit the following information about a selected user:

Type

Enables you to modify the type of user.

Name

Enables you to change the user name of the selected user.

Email

Enables you to edit the email address of the selected user.

Role

Enables you to edit the role that is assigned to the user. This field is displayed only when the selected user is a local user, remote user, or remote group.

Related tasks

[Editing the user settings](#) on page 131

Related references

[Performance Manager user roles and capabilities](#) on page 129

Manage Users page on page 133

Managing user authentication

You can configure Performance Manager to use an authentication server, using LDAP or Active Directory, for authenticating user access to Performance Manager.

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 Performance Manager.

The following LDAP servers are compatible with the management server:

- Microsoft Active Directory
- OpenLDAP
- IBM Lotus LDAP
- Netscape LDAP Server

You can use one of the following predefined authentication services or specify your own authentication service:

- Microsoft Active Directory

Note: 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 Performance 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 Performance 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.

Related tasks

[Enabling remote authentication](#) on page 138

Enabling remote authentication

You can enable remote authentication using OPEN LDAP or Active Directory so that the management server can communicate with your authentication servers, and so users of the authentication servers can use Performance Manager and manage the storage objects and data.

Before you begin

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

About this task

If remote authentication is disabled, remote users or groups cannot access Performance Manager.

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

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > Authentication**.
3. In the **Authentication** dialog box, select **Enable Remote Authentication**.
4. In the Authentication Service field, select either Active Directory or Open LDAP, and then enter the applicable information:

If you are using...	Enter the following information...
Active Directory	<ul style="list-style-type: none">• Authentication server administrator name• Administrator password• Base distinguished name (using the appropriate Active Directory notation)
Open LDAP	<ul style="list-style-type: none">• Bind distinguished name (using appropriate LDAP notation)• Bind password• Base distinguished name

If 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 Performance Manager might reduce the authentication time.

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

Related references

Performance Manager user roles and capabilities on page 129

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 Performance Manager. You might disable nested groups when you want to improve Active Directory authentication response time.

Before you begin

You must be logged in as an Active Directory domain user to perform this task. Logging in as an Active Directory administrator is not required.

About this task

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

Steps

1. Click **Administration > Setup Options**.
2. In the **Setup Options** dialog box, click **Management Server > Authentication**.
3. In the **Authentication Service** field, select **Others**.
4. In the **Member** field, change the member information from member:1.2.840.113556.1.4.1941: to member.
5. Click **Save and Close**.

Setting up authentication services

Authentication services enable the authentication of remote users or groups in an authentication server before providing them access to Performance Manager. You can authenticate users by using

the predefined authentication services, such as Active Directory or OpenLDAP, or by configuring your own authentication mechanism.

Before you begin

- You must have enabled remote authentication.
- You must have the OnCommand Administrator role or the Storage Administrator role.

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > Authentication**.
3. In the **Authentication** dialog box, select one of the following authentication services:

If you select...	Then do this...
Active Directory	<div><div>a. Enter the administrator name and password. You can specify the administrator name in one of the following formats:</div><div><ul style="list-style-type: none">• domainname\username• username@domainname• Bind Distinguished Name, using the appropriate LDAP notation.</div><div>b. Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is dc=ou,dc=domain,dc=com.</div></div>
OpenLDAP	<div><div>a. Enter the bind distinguished name and bind password.</div><div>b. Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is dc=ou,dc=domain,dc=com.</div></div>

If you select...	Then do this...
Others	<ul style="list-style-type: none">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 <code>ou@domain.com</code>, then the base distinguished name is <code>dc=ou,dc=domain,dc=com</code>.c. Specify the LDAP protocol version that is supported by the authentication server.d. Enter the user name, group membership, user group, and member attributes.

Note: If you want to modify the authentication service, ensure that you first delete any existing authentication servers and then add new authentication servers.

4. Click **Save and Close**.

Related tasks

[Enabling remote authentication](#) on page 138

Related references

[Performance Manager user roles and capabilities](#) on page 129

Adding authentication servers

You can add authentication servers and enable remote authentication on the management server to enable remote users within the authentication server to access Performance 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 or the Storage Administrator role.

About this task

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

Steps

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

Result

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

After you finish

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

Related concepts

[*Authentication with Active Directory or OpenLDAP*](#) on page 127

Related tasks

[*Enabling remote authentication*](#) on page 138

[*Setting up authentication services*](#) on page 139

[*Testing the configuration of authentication servers*](#) on page 143

Related references

[*Performance Manager user roles and capabilities*](#) on page 129

Editing authentication servers

You can change the port that Performance Manager uses to communicate with your authentication server. You cannot configure Secure LDAP (LDAPS).

Before you begin

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

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > Authentication**.
3. In the **Authentication** dialog box, in the Servers area, select the authentication server that you want to edit, and then click **Edit**.
4. In the **Edit Authentication Server** dialog box, edit the port details.
5. Click **Save and Close**.

Related references

[Performance Manager user roles and capabilities](#) on page 129

Testing the configuration of authentication servers

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

Before you begin

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

About this task

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

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > Authentication**.
3. In the **Authentication** dialog box, click **Test**.
4. In the **Test User** dialog box, specify the user name and password of the remote user or group, and then click **Test**.

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

Related tasks

[Enabling remote authentication](#) on page 138

[Setting up authentication services](#) on page 139

[Adding authentication servers](#) on page 141

Deleting authentication servers

You can delete an authentication server if you want to prevent Performance Manager 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 or the Storage 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 Performance Manager.

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > Authentication**.

3. In the **Authentication** dialog box, in the Servers area, select one or more authentication servers that you want to delete, and then click **Delete**.
4. Click **Yes** to confirm the delete request.

Related references

[Performance Manager user roles and capabilities](#) on page 129

Page descriptions for user authentication

You can specify how users are authenticated when accessing Performance Manager.

The topics below display when you click **Help** on the appropriate page.

Authentication dialog box

You can use the Authentication dialog box to configure the management server to communicate with your authentication server and authenticate remote users in the authentication server.

Enable Remote Authentication

This area allows you to enable or disable remote authentication. You can enable remote authentication to enable the management server to authenticate remote users within the configured authentication servers.

Authentication Service

Enables you to configure the management server to authenticate users in directory service providers, such as Active Directory, OpenLDAP, or specify your own authentication mechanism. You can specify an authentication service only if you have enabled remote authentication.

Active Directory

- **Administrator Name**
Specifies the administrator name of the authentication server. The name must include the domain name and user name. For example, domain\admin.
- **Password**
Specifies the password to access the authentication server.
- **Base Distinguished Name**
Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **dc=ou,dc=domain,dc=com**.

OpenLDAP

- **Bind Distinguished Name**
Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server.
- **Bind Password**
Specifies the password to access the authentication server.
- **Base Distinguished Name**
Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is `ou@domain.com`, then the base distinguished name is **`dc=ou,dc=domain,dc=com`**.

Others

- **Bind Distinguished Name**
Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server that you have configured.
- **Bind Password**
Specifies the password to access the authentication server.
- **Base Distinguished Name**
Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is `ou@domain.com`, then the base distinguished name is **`dc=ou,dc=domain,dc=com`**.
- **Version**
Specifies the Lightweight Directory Access Protocol (LDAP) version that is supported by your authentication server. You can specify whether the protocol version must be automatically detected or set the version to 2 or 3.
- **User Name Attribute**
Specifies the name of the attribute in the authentication server that contains user login names to be authenticated by the management server.
- **Group Membership Attribute**
Specifies a value that assigns the management server group membership to remote users based on an attribute and value specified in the user's authentication server.
- **UGID**
If the remote users are included as members of a `GroupOfUniqueNames` object in the authentication server, this option enables you to assign the management server group membership to the remote users based on a specified attribute in that `GroupOfUniqueNames` object.
- **Member**

Specifies the attribute name that your authentication server uses to store information about the individual members of a group.

- **User Object Class**
Specifies the object class of all users in the remote authentication server.
- **Group Object Class**
Specifies the object class of all groups in the remote authentication server.

Note: If you want to modify the authentication service, you must first delete any existing authentication servers and add new authentication servers.

Servers

This 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 servers.

Command buttons

Enables you to add, edit, or delete authentication servers.

- **Add**
Displays the Add Server dialog box for adding an authentication server. You specify the name or IP address of the server and the port number.
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**
Displays the Edit Server dialog box for editing the settings for an authentication server. You can edit the name or IP address of the server and the port number.
- **Remove**
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 area enables you to test your configuration.

Test

Validates the configuration of your authentication server by authenticating a remote user or group.

While testing, if you specify only the user name, the management server searches for the remote user in the authentication server, but does not authenticate the user. If you specify both the user name and password, the management server searches and authenticates the remote user.

You cannot test the authentication if remote authentication is disabled.

Related tasks

[*Enabling remote authentication*](#) on page 138

[*Setting up authentication services*](#) on page 139

[*Adding authentication servers*](#) on page 141

[*Testing the configuration of authentication servers*](#) on page 143

[*Deleting authentication servers*](#) on page 144

Managing security certificates

You can configure HTTPS in the Performance 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 Performance Manager is not being intercepted. You can also view the certificate to verify the content of a regenerated certificate or to view alternate URL names from which you can access Performance Manager.

Before you begin

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

Steps

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

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

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

Related tasks

[Adding users](#) on page 129

[Generating an HTTPS security certificate](#) on page 150

[Downloading an HTTPS certificate signing request](#) on page 151

[Installing an HTTPS security certificate](#) on page 152

[Restarting the Performance Manager virtual machine](#) on page 150

Restarting the Performance Manager virtual machine

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

Before you begin

The virtual appliance is powered on.

You are logged in to the maintenance console as the maintenance user of Performance Manager.

About this task

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

Steps

1. Access the maintenance console.
2. Select **System Configuration > Reboot Virtual Machine**.
3. Start the Performance Manager GUI from your browser and log in.

Related tasks

[Adding users](#) on page 129

[Viewing the HTTPS security certificate](#) on page 149

[Generating an HTTPS security certificate](#) on page 150

[Downloading an HTTPS certificate signing request](#) on page 151

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

Note: If connections that enable performance monitoring are currently configured between the Unified Manager server and one or more Performance Manager servers, executing this task

invalidates those connections and deactivates any further performance monitoring updates from Performance Manager servers to the Unified Manager GUI. You must reactivate those connections after completing this task.

Steps

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

Important: You must restart the Performance Manager virtual machine before the new certificate will take effect. This can be done from the System Configuration option in the NetApp maintenance console or from the VM console.

After you finish

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

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

Related tasks

[Adding users](#) on page 129

[Viewing the HTTPS security certificate](#) on page 149

[Downloading an HTTPS certificate signing request](#) on page 151

[Restarting the Performance Manager virtual machine](#) on page 150

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. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > HTTPS**.

3. Click **Download HTTPS Certificate Signing Request**.
4. Save the <hostname>.csr file.

After you finish

You can provide the file to a Certificate Authority to sign and then install the signed certificate.

Related tasks

[Adding users](#) on page 129

[Viewing the HTTPS security certificate](#) on page 149

[Generating an HTTPS security certificate](#) on page 150

[Installing an HTTPS security certificate](#) on page 152

Installing an HTTPS security certificate

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

Before you begin

You must have completed the following actions:

- Downloaded the Certificate Signing Request file and had it signed by a Certificate Authority
- Saved the certificate chain in PEM format
- Included all certificates in the chain, from the server certificate to the root signing certificate

You must have the OnCommand Administrator role.

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > HTTPS**.
3. Click **Install HTTPS Certificate**.
4. In the dialog box that displays, click **Browse** to locate the file to upload.
5. Select the file and 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-----
<Root signing certificate>
-----END CERTIFICATE-----
```

Page descriptions for certificate management

You can use the Configure Settings dialog box to view the current security certificates and to generate new HTTPS certificates.

The topics below display when you click **Help** on the appropriate page.

HTTPS Certificate dialog box

You can use the HTTPS Certificate dialog box to view the current security certificate, download a certificate signing request, generate a new HTTPS certificate, or install a new HTTPS certificate.

You must have the OnCommand Administrator role.

HTTPS Certificate

You can perform the following operations:

View HTTPS Certificate

Enables you to view the current HTTPS certificate. If you have not generated a new HTTPS certificate, this is the certificate that was generated with your installation.

Regenerate HTTPS Certificate

Enables you to generate an HTTPS certificate, which replaces the previous security certificate. The new certificate is in effect after you restart the management server.

Download HTTPS Certificate Signing Request

Downloads a certification request for the currently installed HTTPS certificate. Your browser prompts you to save the <hostname>.csr file so that you can provide the file to a Certificate Authority to sign.

Install HTTPS Certificate

Enables you to upload and install a security certificate after a Certificate Authority has signed and returned it. The new certificate is in effect after you restart the management server.

Related tasks

[Viewing the HTTPS security certificate](#) on page 149

[Generating an HTTPS security certificate](#) on page 150

[Downloading an HTTPS certificate signing request](#) on page 151

[Installing an HTTPS security certificate](#) on page 152

[Working with HTTPS security certificates](#) on page 18

Managing event notification

You can set up an SMTP server to enable email communication from Performance Manager and configure email alerts. The email alerts notify you about events on the cluster.

Configuring email settings

You can configure SMTP settings for the Performance Manager server to send email notifications when an event is generated. You can specify the corresponding mail server to be used.

Before you begin

The following information must be available:

- Email address from which the alert notification is sent
- Host name (or IP address), user name, password, and default port to configure the SMTP server

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

Steps

1. Click **Administration > Configure Settings**.
2. In the **Configure Settings** dialog box, click **Management Server > Email**.
3. In the **Email** dialog box, configure the appropriate settings.

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.

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

The user name and password are only required if SMTP authorization is enabled.

4. Click **Test** to confirm whether recipients receive email alerts using the SMTP settings.

Related tasks

[Configuring email alerts](#) on page 156

Related references

[Performance Manager user roles and capabilities](#) on page 129

Configuring email alerts

You can specify which incidents from Performance Manager to alert on and the email recipients for those alerts. You can receive alerts for all new incidents, disable all email alerts, or exclude email alerts caused by a QoS policy group limit. By default, alerts are sent for all new incidents.

Before you begin

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

Steps

1. Click **Administration > Manage Alerts**.
2. In the **Configure Email Alerts** dialog box, configure the appropriate settings.

Note: For Email Recipients, use a comma or semicolon, with or without spaces, to separate the addresses. If you enter several addresses, such as by copying and pasting from an email client, the addresses are automatically separated with commas after you click **Save**.

Related tasks

[Configuring email settings](#) on page 155

Related references

[Performance Manager user roles and capabilities](#) on page 129

Page descriptions for notification management

You can manage event notifications, such as setting up an SMTP server and configuring email alerts, to have Performance Manager notify you about various cluster events.

The topics below display when you click **Help** on the appropriate page.

Email dialog box

You can configure an SMTP server that the Performance Manager server uses to send email notifications when an event is generated. You can also specify a From address.

This dialog box enables you to configure the following SMTP server settings:

From Address

Specifies the address that recipients will see in the From field of their email client.

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. If SMTP authentication is not enabled on the SMTP server, this field is optional.

Password

Specifies the SMTP password. If SMTP authentication is not enabled on the SMTP server, this field is optional.

Port

Specifies the port that is used by the SMTP host server to send alert notification.

The default value is 25.

Use STARTTLS

A mechanism to provide secure communication 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.

Related tasks

[Configuring email settings](#) on page 155

Related references

[Performance Manager user roles and capabilities](#) on page 129

Configure Email Alerts dialog box

You can specify which incidents from Performance Manager to alert on and the email recipients for those alerts. You can also disable all email alerts for all recipients. Email alerts are sent immediately after an incident is detected.

The following options are displayed:

Send For

This section lets you select to have email alerts sent for all incidents or to exclude incidents caused by a QoS policy group limit, when workloads have exceeded the throughput limit. You can also disable all email alerts.

Send To

This section lets you type the address of each email recipient. To remove a recipient, you can delete the appropriate address.

Note: For Email Recipients, use a comma or semicolon, with or without spaces, to separate the addresses. If you enter several addresses, such as by copying and pasting from an email client, the addresses are automatically separated with commas after you click **Save**.

Related tasks

[*Configuring email alerts*](#) on page 156

Related references

[*Performance Manager user roles and capabilities*](#) on page 129

Troubleshooting common issues

There are common issues that you might encounter when using Performance Manager. You can take corrective actions to resolve these issues on your own.

Unknown authentication error

Issue

When you are performing an authentication-related operation, such as adding, editing, deleting, or testing remote users or groups, the following error message might be displayed: Unknown authentication error.

Cause

This problem can occur if you have set an incorrect value for the following:

- Administrator Name of the Active Directory authentication service
- Bind Distinguished Name of the OpenLDAP authentication service

Corrective action

1. Click **Administration > Configure Settings**
2. In the Configure Settings dialog box, click **Management Server > Authentication**.
3. Based on the authentication service that you have selected, enter the appropriate information for Administrator Name or Bind Distinguished Name in the Authentication dialog box.
4. Click **Save and Close**.

Icons are misaligned in Internet Explorer

Issue

Icons and text are misaligned when you use Internet Explorer.

Cause

This problem can occur if you are using Internet Explorer in Compatibility View, which is not a supported browser setting.

Corrective action

1. Press F12 to open Internet Explorer Developer Tools.

2. Select **Browser Mode** from the toolbar to display the browser version used to open the application.
3. Select **Document Mode** from the toolbar and select the Standards mode of the browser version used to open the application.
For example, if you are using Internet Explorer 9 to open the application, select **Browser Mode > Internet Explorer 9**, and then select **Document Mode > Internet Explorer 9 Standards**.

LDAP server slow to respond

Issue

The LDAP server takes a long time to respond to queries.

Cause

Supporting nested groups causes the LDAP server to slow down.

Corrective action

If you use Active Directory, you can speed authentication by disabling support for nested groups in Performance Manager. However, if you choose to disable nested groups, you must ensure that users are direct members of the groups added to Performance Manager.

To disable nested group support, follow these steps:

1. Click **Administration > Configure Settings > Authentication** to display the Authentication dialog box.
2. Select the **Enable Remote Authentication** check box.
3. In the **Authentication Service** drop-down menu, select **Others**.
4. In the **Member** box, type “member”.
5. Click **Save and Close**.

Issue with adding LDAP using Other authentication services

Issue

When you select Other in the Authentication dialog box, the user and groupObjectClass retain the values from the previously selected template. If the LDAP server does not use the same values, the operation might fail.

Cause

The users are not configured correctly in OpenLDAP.

Corrective action

You can manually fix this issue using one of the following workarounds.

If your LDAP user and object classes are user and group, respectively, then perform the following steps:

1. Click **Administration > Configure Authentication Settings** to display the Authentication dialog box.
2. In the **Authentication Service** drop-down menu, select **Active Directory** and then select **Others**.
3. Complete the text fields.

If your LDAP user and group object classes are posixAccount and posixGroup, respectively, then perform the following steps:

1. Click **Administration > Configure Authentication Settings** to display the Authentication dialog box.
2. In the **Authentication Service** drop-down menu, select **OpenLDAP** and then select **Others**.
3. Complete the text fields.

Copyright information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- access checks
 - introduction to using RBAC to enable application administrator [126](#)
- access roles (RBAC)
 - See* RBAC
- Active Directory
 - setting up authentication services [139](#)
 - using to enable remote authentication [138](#)
- Add Cluster dialog box [123](#)
- Add User dialog box [134](#)
- adding
 - authentication servers [141](#)
 - clusters [117](#)
 - remote groups [129](#)
 - remote users [129](#)
- adding clusters
 - requirements for [118](#)
- administrator roles
 - See* RBAC
- administrators
 - introduction to using RBAC to restrict functionality access to selected [126](#)
 - OnCommand [128](#)
 - storage [128](#)
- aggregate overcommitment
 - considerations for using with FlexVol volumes [34](#)
- aggregates
 - availability [43](#)
 - capacity states [44](#)
 - states [43](#)
- alerts
 - configuring [20](#), [155](#), [156](#)
 - configuring your environment to send [13](#)
- analysis, performance
 - how the expected range is used in [51](#)
- analyzing
 - component contention [108](#)
 - performance incidents [108](#)
 - workload response time [53](#)
- architecture
 - MetroCluster [38](#)
- assigning
 - user roles [129](#)
- authentication
 - Active Directory [127](#), [137](#)

- adding servers [141](#)
- deleting servers [144](#)
- editing servers [143](#)
- enabling remote [138](#)
- OpenLDAP [127](#), [137](#)
- servers [145](#)
- testing for remote users and groups [143](#)
- troubleshooting unknown authentication error [159](#)
- authentication services
 - setting up using Active Directory [139](#)
 - setting up using OpenLDAP [139](#)
- authentication, remote
 - disabling nested groups [139](#)
- AutoSupport
 - enabling [21](#)
 - enabling periodic messages [14](#)
 - generating [21](#)
 - options offered by the Performance Manager setup wizard [10](#)
 - sending [21](#)
 - sending an on-demand message [16](#)
 - viewing description [21](#)
 - what it does [15](#)

B

- benefits
 - of using SVMs [31](#)
- browsers
 - requirements [70](#)
- bully workloads
 - identifying as part of a performance incident [89](#)

C

- cache hit ratio
 - analyzing [79](#)
- cache storage
 - ratio of reads and writes for a workload [84](#)
- capabilities
 - table of roles associated with [129](#)
- capacity states [44](#)
- certificates
 - downloading HTTPS certificate signing requests [151](#)
 - generating HTTPS security certificates [150](#)
 - installing HTTPS security certificates [152](#)

- security, downloading signing request [153](#)
- security, regenerating [153](#)
- security, viewing [153](#)
- viewing HTTPS security certificates [149](#)
- charts
 - Contention History [114](#)
- Chrome
 - browser requirements [70](#)
- client software
 - supported versions [70](#)
- cluster
 - components [24](#)
 - concepts [24](#)
- cluster components
 - causes of contention [63](#)
 - performance trending [78](#)
- cluster components in contention
 - identifying contributing bully workloads [89](#)
- cluster configuration changes
 - monitoring [66](#)
- cluster network [26](#)
- cluster operations
 - analyzing in Performance Manager [54](#)
- clusters
 - adding [117](#)
 - checking the health of [95](#)
 - configuring [117](#), [122](#)
 - deleting [120](#)
 - description of [24](#)
 - editing settings [120](#)
 - how discovery process works [118](#)
 - options to add offered by the Performance Manager setup wizard [10](#)
 - removing [120](#)
 - requirements for adding [118](#)
 - viewing inventory list [119](#)
- clusters and HA pairs [26](#)
- comments
 - how to send feedback about documentation [164](#)
- component contention
 - analyzing [108](#)
- components
 - of a system running clustered Data ONTAP [24](#)
- concepts
 - Performance Manager [24](#)
- configurations, MetroCluster
 - checking the health of partner clusters [95](#)
- Configure Email Alerts dialog box [157](#)
- Configure Email Settings dialog box [156](#)
- Configure Network Settings dialog box

- content and purpose [22](#)
- configuring
 - alerts [155](#), [156](#)
 - authentication [145](#)
 - clusters [117](#), [122](#)
 - data sources [117](#), [122](#)
 - DNS [17](#)
 - email notification settings [155](#)
 - HTTPS settings [153](#)
 - initial settings after installation [10](#)
 - network settings [17](#), [20](#)
 - notifications [20](#), [155](#), [156](#)
 - NTP server [21](#)
 - security certificates [149](#), [153](#)
 - user authentication [137](#), [145](#)
 - users [133](#), [137](#), [145](#)
 - your environment [11](#)
- contention
 - analyzing [108](#)
 - causes of [63](#)
 - definition [63](#)
 - identifying [63](#)
- Contention History charts
 - types of, described [114](#)
- contention, cluster component
 - identifying contributing bully workloads [89](#)
- copying link to Volume Details page [73](#)
- creating
 - database users [129](#)
 - local users [129](#)

D

- Dashboard
 - description of [105](#)
 - Incidents section [107](#)
 - Quick Takes section [105](#)
- data breakdown charts [84](#)
- data collection
 - using to monitor workload performance [46](#)
- data network [26](#)
- data processing
 - definition [37](#)
- data sources
 - configuring [117](#), [122](#)
- database users
 - creating [129](#)
 - defined [128](#)
- deleting
 - authentication servers [144](#)

- clusters [120](#)
- users [132](#)
- DHCP
 - enabling [17](#)
- dialog boxes
 - Configure Network Settings [22](#)
- disaster recovery group
 - MetroCluster [38](#)
- discovery
 - of clusters [118](#)
- disk failures
 - responding to incidents [101](#)
- displaying
 - the Incident Details page [87](#)
- DNS
 - configuring [17](#)
- documentation
 - how to send feedback about [164](#)
 - list of [9](#)

E

- Edit Cluster dialog box [124](#)
- Edit User dialog box [135](#)
- editing
 - authentication servers [143](#)
 - cluster settings [120](#)
 - network settings [17](#)
 - user settings [131](#)
- email alerts
 - configuring [20](#), [155](#), [156](#)
- email notification settings
 - configuring [155](#)
- email notifications
 - configuring [156](#)
 - disabling [156](#)
 - postponing [156](#)
- enabling
 - AutoSupport [21](#)
 - DHCP [17](#)
 - periodic AutoSupport [14](#)
- environment
 - setup [11](#)
- events
 - caused by cluster configuration changes [66](#)
 - configuring notifications for [156](#)
 - configuring your environment to send notifications about [13](#)
 - impacting storage performance [58](#)
 - performance incidents [59](#)

- types of [58](#)
- expected range
 - analysis [48](#)
 - collection interval [46](#)
 - definition of [50](#)
 - how it is used in performance analysis [51](#)
 - measurements [48](#)
 - updating [46](#)
- expected values [48](#)
- exporting data to CSV [72](#)

F

- feedback
 - how to send comments about documentation [164](#)
- Firefox
 - browser requirements [70](#)
- FlexVol volumes
 - considerations for using thin provisioning with [34](#)
 - controlling I/O performance [36](#)
 - how moving them works [81](#)
 - with SVMs, explained [29](#)
- footprint
 - volume, described [33](#)

G

- generating
 - AutoSupport [21](#)
- graphs of performance data [70](#)
- groups
 - introduction to using RBAC to define user roles for managing [126](#)
 - testing remote authentication [143](#)
- groups, nested
 - disabling remote authentication of [139](#)
- GUI
 - logging in [69](#)

H

- HA configurations
 - definition of [25](#)
- HA pairs and clusters [26](#)
- hardware
 - components described [25](#)
 - HA components described [25](#)
- history, performance
 - information displayed in Contention History charts [114](#)

host names

changing Performance Manager [12](#)

HTTPS

configuring [149](#), [153](#)

downloading certificate signing requests [151](#)

generating new security certificates [150](#)

installing security certificates [152](#)

viewing the security certificate [149](#)

HTTPS certificates

downloading a new [18](#)

installing a new [18](#)

regenerating [18](#)

viewing [18](#)

working with [18](#)

I

icons misaligned troubleshooting [159](#)

Incident Details page

Activity History chart [114](#)

displaying [87](#)

Incident History chart [114](#)

Top Victim Workloads chart [114](#)

Total Write Throughput chart [114](#)

Workload Details table explained [112](#)

Workloads and cluster components involved [109](#)

incidents

caused by disk failure [101](#)

identifying shark workloads involved in [91](#)

investigating [101](#)

remediations [116](#)

resolving [116](#)

searching for [68](#)

viewing workloads and components involved [109](#)

incidents, performance

analysis of [59](#)

analyzing those caused by volume moves [79](#)

caused by HA takeover [103](#)

caused by policy group throttling [99](#)

checking the health of partner clusters after [95](#)

definition of [59](#)

determining impact to workload performance [62](#)

identifying the bully workloads [89](#)

investigating [103](#)

investigating slow response time [76](#)

notification of [59](#)

states [59](#)

viewing [87](#)

Infinite Volumes

with SVMs, explained [29](#)

information

how to send feedback about improving
documentation [164](#)

Internet Explorer

browser requirements [70](#)

IOPS performance

introduction to using Storage QoS to monitor [34](#)

issues, performance

investigating [75](#)

L

LDAP server slow to respond

troubleshooting [160](#)

LDAP user of OpenLDAP server

troubleshooting [160](#)

limits

introduction to using Storage QoS to manage
workload throughput [34](#)

local users

creating [129](#)

defined [128](#)

logging in [69](#)

logical storage

defined [29](#)

M

Macintosh

supported versions [70](#)

maintenance console

restarting the Performance Manager virtual machine
[150](#)

restarting the virtual machine [150](#)

maintenance user

defined [128](#)

what it does [126](#)

Manage Data Sources page

explained [122](#)

Manage users page [133](#)

managing

users [126](#), [133](#)

MBps performance

introduction to using Storage QoS to monitor [34](#)

messages

sending on-demand AutoSupport [16](#)

MetroCluster

architecture [38](#)

disaster recovery group [38](#)

illustration

- MetroCluster illustration [38](#)
- MetroCluster configurations
 - analyzing performance incidents for [92, 98](#)
 - checking the health of partner clusters [95](#)
 - monitoring performance of [55](#)
 - monitoring volumes during switchover and switchback [56](#)
 - SVM data protection in [41](#)
- misalignment of icons troubleshooting [159](#)
- modifying
 - user settings [131](#)
- monitoring
 - cluster configuration changes [66](#)
 - MetroCluster configurations [55](#)
 - volumes in MetroCluster configurations [56](#)
- moved volumes
 - analyzing the performance impact from [79](#)

N

- names
 - changing Performance Manager host [12](#)
- navigating the GUI [68](#)
- nested groups
 - disabling remote authentication of [139](#)
 - disabling support for to speed performance [160](#)
- network processing
 - definition [37](#)
- network settings
 - changing the Performance Manager host name [12](#)
 - configuration values in Configure Network Setting dialog box [22](#)
 - editing [17](#)
- Network Time Protocol
 - See* NTP
- nodes
 - description of [25](#)
 - single-node cluster
 - See* single-node clusters
- notification
 - configuring email settings [155](#)
- notifications
 - configuring [20, 155, 156](#)
- notifications for events
 - configuring [156](#)
- notifications, alert
 - configuring your environment to send [13](#)
- NTP server
 - defining [21](#)
- NTP settings

- configuring [17](#)

O

- on-demand AutoSupport messages
 - sending [16](#)
- OnCommand administrators
 - defined [128](#)
- OnCommand maintenance console
 - role of maintenance user [126](#)
- Open LDAP
 - using to enable remote authentication [138](#)
- OpenLDAP
 - setting up authentication services [139](#)
- operations
 - read and write requests [84](#)
- operators
 - defined [128](#)

P

- page descriptions for
 - alert notifications [156](#)
 - AutoSupport [20](#)
 - cluster management [122](#)
 - network settings [20](#)
 - NTP settings [20](#)
 - security certificates [153](#)
 - system setup [20](#)
 - user authentication [145](#)
 - user management [133](#)
- passwords
 - changing [132](#)
- peak deviation [48](#)
- peak values [48](#)
- performance
 - expected range [50](#)
 - storage events in Performance Manager [58](#)
- performance analysis
 - how the expected range is used in [51](#)
- performance events
 - types of [58](#)
- performance incidents
 - analysis of [59](#)
 - analyzing [108](#)
 - analyzing for MetroCluster configurations [92](#)
 - analyzing in a MetroCluster configuration [98](#)
 - analyzing those caused by volume moves [79](#)
 - caused by disk failure [101](#)
 - caused by HA takeover [103](#)

- caused by MetroCluster resource contention [55](#)
- caused by policy group throttling [99](#)
- checking the health of partner clusters after [95](#)
- definition of [59](#)
- detection [53](#)
- identifying bully workloads [89](#)
- identifying shark workloads involved in [91](#)
- identifying victim workloads [88](#)
- in all flash aggregates [59](#)
- in Flash Pool aggregates [59](#)
- in HDD aggregates [59](#)
- investigating [101](#), [103](#)
- investigating slow response time [76](#)
- list of [107](#)
- notification of [59](#)
- resolving [108](#)
- states [59](#)
- viewing [87](#)
- performance issues
 - finding suggested actions for [116](#)
 - investigating [75](#)
- Performance Manager
 - analysis of performance impact [62](#)
 - changing the host name [12](#)
 - events [58](#)
 - features [8](#)
 - graphs [70](#)
 - how it uses cluster operations [54](#)
 - how it uses volume response time [53](#)
 - how the expected range is used in analysis [51](#)
 - incidents [59](#)
 - introduction [8](#)
 - navigating the interface [70](#)
 - troubleshooting common issues [159](#)
 - types of workloads monitored [47](#)
- performance measurements [48](#)
- performance statistics, historical
 - displayed in Contention History charts [114](#)
- performance threshold [48](#)
- performance, workload
 - collecting data to monitor [46](#)
- periodic support messages
 - enabling [14](#)
- physical storage
 - adding clusters [117](#)
 - editing cluster settings [120](#)
 - Manage Data Sources page explained [122](#)
 - removing clusters [120](#)
- policy groups
 - analyzing [99](#)

- how maximum throughput works [35](#)
 - impact of throttling on response time [84](#)
- port requirements
 - Performance Manager setup [19](#)
- ports
 - editing, for authentication servers [143](#)
- printing a page [74](#)
- product documentation
 - list of [9](#)
- protocol requirements
 - Performance Manager setup [19](#)

Q

- QoS policy group throttling
 - responding to performance incidents caused by [99](#)

R

- range, expected
 - how it is used in performance analysis [51](#)
- RBAC
 - definition [126](#)
 - introduction to managing groups of users by using [126](#)
- remote authentication
 - disabling nested groups [139](#)
 - enabling [138](#)
 - servers [145](#)
- remote groups
 - adding [129](#)
 - defined [128](#)
 - testing authentication [143](#)
- remote users
 - adding [129](#)
 - defined [128](#)
 - testing authentication [143](#)
- removing
 - clusters [120](#)
- requirements
 - browser [70](#)
 - Performance Manager protocol and port setup [19](#)
- response time
 - analysis of workload performance [53](#)
 - analyzing volume performance [53](#)
 - cluster components [84](#)
 - detecting performance incidents [53](#)
 - disk operations for a workload [84](#)
 - impact from policy group throttling [84](#)
 - read and write requests [84](#)

- total CPU time used by a workload [84](#)
- trending [78](#)
- viewing impact from performance incidents [62](#)
- viewing workloads with highest deviation of [88](#)
- response time, slow
 - investigating [76](#)
- role-based access control
 - See* RBAC
- roles
 - assigning to users [129](#)
 - defined [128](#)
 - table of capabilities associated with [129](#)

S

- search bar
 - using to find storage objects [121](#)
- searching
 - for storage objects [121](#)
- searching for objects [68](#)
- security certificates
 - configuring [149](#), [153](#)
 - downloading a new [18](#)
 - downloading HTTPS certificate signing requests [151](#)
 - downloading signing request [153](#)
 - generating, HTTPS [150](#)
 - installing a new [18](#)
 - installing, HTTPS [152](#)
 - regenerating [18](#), [153](#)
 - viewing [18](#), [153](#)
 - viewing, HTTPS [149](#)
 - working with [18](#)
- sending
 - AutoSupport [21](#)
- servers
 - required ports for Performance Manager setup [19](#)
- setting up
 - email notification settings [155](#)
 - SMTP server email notifications [155](#)
- settings, network
 - configuration values in Configure Network Setting dialog box [22](#)
- settings, NTP
 - configuring [17](#)
- setup
 - post-deployment [11](#)
- sharing data
 - as a CSV file [72](#)
 - copying a page link [73](#)
 - printing a page [74](#)

- shark workloads
 - identifying in performance incidents [91](#)
- sharks defined [65](#)
- single-node clusters
 - description of [24](#)
- slow response time
 - investigating [76](#)
- states
 - for volumes [33](#)
- storage administrators
 - defined [128](#)
- storage objects
 - searching for [121](#)
- Storage QoS
 - effect on non-throttled workloads [35](#)
 - how maximum throughput works [35](#)
 - introduction to managing workload performance by using [34](#)
- Suggested Actions section
 - Incident Details page [116](#)
- suggestions
 - how to send feedback about documentation [164](#)
- SVMs
 - benefits of using [31](#)
 - data protection in MetroCluster configurations [41](#)
 - introduction to using Storage QoS to limit workload throughput to LUNs within [34](#)
 - with FlexVol volumes, explained [29](#)
 - with Infinite Volume, explained [29](#)
- SVMs with FlexVol volumes
 - explained [29](#)
- SVMs with Infinite Volume
 - explained [29](#)
- switchback
 - monitoring volumes during MetroCluster [56](#)
- switchover
 - monitoring volumes during MetroCluster [56](#)
- system running clustered Data ONTAP
 - adding [117](#)
 - removing [120](#)
- system-defined workloads
 - definition [47](#)

T

- testing
 - authentication for remote users and groups [143](#)
- thin provisioning
 - considerations for using with FlexVol volumes [34](#)
- throttling, QoS policy group

- responding to performance incidents caused by [99](#)
- throughput
 - ratio of reads and writes to cache [84](#)
- throughput limits
 - introduction to using Storage QoS to manage workload [34](#)
- time synchronization
 - using NTP settings configuration [17](#)
- troubleshooting
 - adding LDAP user of OpenLDAP server [160](#)
 - LDAP server slow to respond [160](#)
 - misalignment of icons [159](#)
 - Performance Manager [159](#)
 - unknown authentication error [159](#)
- two-node clusters
 - description of [24](#)
- types
 - of users [128](#)
- types of users
 - maintenance user [126](#)

U

- unknown authentication error troubleshooting [159](#)
- user authentication
 - configuring [137](#), [145](#)
- user management [126](#)
- user roles
 - assigning [129](#)
- user-defined workloads
 - definition [47](#)
- users
 - adding [129](#)
 - capabilities associated with [129](#)
 - changing passwords [132](#)
 - configuring [133](#), [137](#), [145](#)
 - creating [129](#)
 - deleting [132](#)
 - editing settings [131](#)
 - introduction to using RBAC to manage groups of [126](#)
 - maintenance user [126](#)
 - managing [133](#)
 - modifying settings [131](#)
 - roles [128](#)
 - testing remote authentication [143](#)
 - troubleshooting LDAP user of OpenLDAP server [160](#)
 - types [128](#)
 - viewing [131](#), [133](#)

- Users and Roles capability
 - See* RBAC

V

- victim workloads [65](#)
- viewing
 - AutoSupport description [21](#)
 - clusters list [119](#)
 - users [131](#)
 - users list [133](#)
- virtual appliances
 - changing the Performance Manager host name [12](#)
- virtual machine
 - restarting [150](#)
- virtual machines
 - required ports for Performance Manager setup [19](#)
- Volume Details page
 - content of [82](#)
 - data breakdown charts [84](#)
 - navigating to [82](#)
- volume footprint
 - described [33](#)
- volume moves
 - analyzing the performance impact of [79](#)
- volume performance
 - analyzing response time [53](#)
- volume workloads
 - performance information [82](#)
- volumes
 - behavior during MetroCluster switchover and switchback [56](#)
 - considerations for using thin provisioning with FlexVol [34](#)
 - defined [32](#)
 - FlexVol, how moving them works [81](#)
 - how moving FlexVol volumes works [81](#)
 - how they work [32](#)
 - introduction to using Storage QoS to limit workload throughput to [34](#)
 - searching for [68](#)
 - states [33](#)
 - SVMs with FlexVol, explained [29](#)
 - SVMs with Infinite, explained [29](#)
- Vservers
 - See* SVMs

W

- Windows

- supported versions [70](#)
- Workload Details table
 - what it displays [112](#)
- workload performance
 - collecting data to monitor [46](#)
- workload response time
 - expected range
 - response time
 - expected range [50](#)
- workloads
 - bullies defined [65](#)
 - effect of throttling on non-throttled workloads [35](#)
 - introduction to using Storage QoS to manage performance [34](#)
 - investigating performance issues [75](#)
 - involved in an incident [109](#)
 - performance details [84](#)
 - performance measurements [48](#)
 - searching for [68](#)
 - system-defined [54](#)
 - types monitored by Performance Manager [47](#)
 - types not monitored by Performance Manager [47](#)
 - user-defined [54](#)
 - victims defined [65](#)
 - viewing impact from performance incidents [62](#)
 - volume details [82](#)
- Workloads involved in incidents
 - displayed in the Workload Details table [112](#)
- workloads, bully
 - identifying as part of a performance incident [89](#)