



OnCommand[®] Insight 6.4

OnCommand Insight Performance 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) 4-NETAPP
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-07593_A0
March 2013

Contents

OnCommand Insight Performance features	6
What you can do with the OnCommand Insight Performance features	6
OnCommand Insight product portfolio	7
Where to find more information about OnCommand Insight	8
OnCommand Insight on the web	8
OnCommand Insight product documentation	9
Videos for OnCommand Insight	10
Setting and managing switch thresholds and performance alerts	11
Adjusting switch thresholds	11
Switch threshold types and formulas	12
Changing or disabling switch thresholds	14
Managing performance alerts	14
Reviewing and confirming alerts	15
Examining performance	16
Searching for performance data	16
Performance trend analysis	18
Identifying reasons for performance changes	20
Monitoring system changes that impact performance	21
Monitoring Fibre Channel switch performance	23
Monitoring data store performance	24
Monitoring storage performance	25
Customizing global general policies	27
General policy types and their hierarchies	29
Identifying performance trends	30
Exporting performance data as an image or file	32
Grouping data in views	32
Monitoring virtual machine performance	34
Viewing virtual machine performance and utilization	34
Setting virtual machine thresholds	35
Determining where to add a data store to achieve load balancing	36
Determining virtual machine contention	37
Selecting an active path for a virtual machine data store	39

Adding a host to a virtual cluster to share a data store	40
Determining if a VM is affecting host performance	41
Troubleshooting performance problems	43
Troubleshooting slow client computers on your network	43
Troubleshooting poor application performance	44
Analyzing traffic flow and identifying congestion	46
Analyzing traffic flow	46
Identifying a congested host	48
Identifying the cause of the congested host	48
Analyzing congested storage ports	49
Analyzing virtual storage performance	50
Balancing resources	53
Analyzing port balance violations	53
Identifying out-of-balance network traffic across host ports	55
Identifying resources for storage allocation	56
Analyzing storage port balance violations	56
Identifying an available storage array	57
Selecting storage ports	57
Identifying volumes with potential capacity	57
Identifying Tier 1 to Tier 2 data migration candidates	58
Balancing data flow across storage array ports	59
Balancing data flow across ISLs	59
Allocating capacity to a new host	60
Allocating capacity to an existing host	62
Performance reports	64
Gathering performance information	64
Reporting array performance for orphaned volumes and tiers	65
Performance reference	67
Automatic storage tiering in Inventory and Performance views	67
Analyze Contention dialog box	68
Analyze dialog box	70
Daily Performance Pattern histogram	71
Data Store Summary tab	72
Disks tab	73
Host Summary tab	75
Internal Volume Summary tab	76

Storage Pool Summary tab	77
Virtual Machine Summary tab	78
Volume Summary tab	79
Application Performance view	81
Backend Volume Performance view	83
Datastore Performance view	86
Disk Performance view	89
Histogram view	92
Host Performance view	93
Internal Volume Performance view	97
Internal Volume Usage of Disk view	100
Performance Chart	101
Performance Distribution view	103
Port Performance Distribution view	104
Settings for Host Virtualization policies	105
Settings for VM Hosts Thresholds	105
Storage Performance view	106
Storage Pool Performance view	110
Storage Resource Performance chart	113
Switch Port Performance view	114
Switch Threshold settings	117
Switch threshold types and formulas	118
Topology view	119
VMDK Performance view	120
VM Distribution view	122
Virtual Internal Volume Performance view	123
Virtual Machine Performance view	126
Virtual Volume Usage view	128
Volume Performance view	130
Volume Usage of Disk view	133
Copyright information	136
Trademark information	137
How to send your comments	138
Index	139

OnCommand Insight Performance features

OnCommand Insight extracts near real-time performance information from your storage environment and maps it to applications, hosts, and access paths. You can search for objects with problematic performance and analyze the related data. These features allow you to use your resources efficiently and help you solve performance problems quickly.

OnCommand Insight Performance includes the following features:

- Heterogeneous Fibre Channel switch fabric performance data for hosts and storage ports
- Automatic discovery, collection, and maintenance of performance metrics
- Search facility linked to performance analysis
- Data about volume congestion and excessive utilization
- Service level performance monitoring

OnCommand Insight provides visibility into the actual storage capacity consumption on every FC port in the monitored environment (host, switch, or storage). You can also obtain statistics on volume, internal volume, and disk performance.

Armed with this information, you can identify expensive tier 1 arrays that have applications that can be moved to tier 2 based on low traffic. After the migration, you can validate that you have not negatively impacted the performance of the applications while balancing your storage usage.

Related concepts

[*Analyzing traffic flow and identifying congestion*](#) on page 46

Related tasks

[*Searching for performance data*](#) on page 16

[*Troubleshooting slow client computers on your network*](#) on page 43

[*Troubleshooting poor application performance*](#) on page 44

What you can do with the OnCommand Insight Performance features

To help storage administrators align SANs for maximum application performance at the lowest cost, OnCommand Insight leverages information about storage resources identified during data acquisition and combines it with real-time performance data. Storage administrators use OnCommand Insight to identify underutilized resources, maintain storage tiers, and troubleshoot performance bottlenecks.

Using the OnCommand Insight **Performance** features, you can do the following:

- Search resources to identify the sources of performance problems quickly.

- Discover critical storage performance metrics and their impact on applications.
- Shift application load to lower-cost tiers.
- Identify the best array and switch resources to allocate to new applications.
- Identify the busiest arrays and find arrays that are not being used efficiently.
- Correlate traffic balance problems to application performance and multipath configuration errors.
- Proactively load-balance the entire application portfolio across the storage fabric to defer capital expenditures.
- Uncover switch-related issues that cause application performance problems.
- Identify array bottlenecks in the context of impacted hosts and applications.
- Analyze the efficiency of virtualized storage.

OnCommand Insight product portfolio

OnCommand Insight is a suite of products that boost the quality and efficiency of your storage environment by providing near real-time visibility, multi-vendor and multi-protocol support, and end-to-end service-level views of your application storage path.

The OnCommand Insight portfolio provides the tools for you to manage your environment. You can do the following:

- Optimize the assets that you have on the floor, according to space, performance, and usage criteria.
- Optimize the business activities, such as application or storage service, on your assets.

The OnCommand Insight portfolio includes the following products:

Assure	Perform	Plan
<ul style="list-style-type: none"> • Enable config SLO • Identify cause of service issues • Plan and validate service changes • Audit changes 	<ul style="list-style-type: none"> • Manage and optimize resource usage • Get storage service performance metrics • Align service tiers 	<ul style="list-style-type: none"> • Manage and plan capacity • Trend, forecast, report • Be cost aware • Enable chargeback/accountability
Availability	Optimize	Efficiency

OnCommand Insight Assure

Helps administrators focus on maintaining storage service on their storage resources. Lets administrators establish global, application, host, or path-based service policies and manage the effects of change, risk, and service-level violations of those policies on availability and performance.

OnCommand Insight Perform	Helps administrators focus on their storage resources. Collects and maintains service performance information to help administrators reclaim under-utilized storage resources, manage tiers, analyze storage virtualization efficiency, improve load balancing, troubleshoot congestion and contention, and improve resource utilization to optimize services.
OnCommand Insight Plan	Helps administrators focus on both storage resources and the services on those assets. Provides administrators with an end-to-end view of resource order, the storage allocation process, and trend usage so that they can manage resource reservations and forecast consumption and resource needs. Administrators can report on costs for storage services by business entity or tier and use this information for chargebacks on storage usage. Business level roll-up reporting helps storage administrators improve capacity planning and optimize consumption by application and tier.

OnCommand Insight Assure, OnCommand Insight Perform, and OnCommand Insight Plan use the same integrated discovery component. This discovery component provides visibility into storage availability, resource inventory, and host-to-storage access paths. Administrators can also gain visibility into the path from virtual machines to arrays for organizations deploying VMware ESX technology combined with VMware vCenter (VirtualCenter) technology.

The Inventory features are the foundation features for the OnCommand Insight suite and are used in OnCommand Insight Assure, OnCommand Insight Perform, and OnCommand Insight Plan.

OnCommand Insight Assure, OnCommand Insight Perform, and OnCommand Insight Plan also use the same report authoring tools and the same Data Warehouse platform.

Where to find more information about OnCommand Insight

You can find more information about OnCommand Insight on the NetApp Support Site and in other OnCommand Insight documentation.

OnCommand Insight on the web

For comprehensive, up-to-date information about OnCommand Insight, use these NetApp web site resources.

- OnCommand Insight product web site at www.netapp.com/oncommandinsight
- The NetApp Support Site at: support.netapp.com
- The OnCommand Insight data source Interoperability Matrix at support.netapp.com/NOW/products/interoperability.

OnCommand Insight product documentation

The OnCommand Insight products are accompanied by a set of guides.

You can access the following documents from the NetApp Support Site at support.netapp.com/documentation/productsatoz/index.html.

OnCommand Insight Getting Started Guide Helps new OnCommand Insight users set up and customize their installed system and begin using it for improved efficiency.

Note: Information about OnCommand Insight is available from the OnCommand Insight Client Help menu and from the OnCommand Insight Administration portal.

OnCommand Report Installation and Administration Guide Provides an overview of suite architecture with instructions to get the system up and running, discover the logical and physical storage resources in your storage environment, and perform administrative tasks.

Describes the configuration parameters and some installation procedures for data sources used by OnCommand Insight. Provides recommended methods of discovering your storage environment for the OnCommand Insight deployment.

OnCommand Insight Inventory User Guide Provides information about the tools and features to help you gain visibility into your SAN or NAS inventory environment. The Inventory features are the foundation used by all other products in the OnCommand Insight suite. After looking at the *OnCommand Insight Getting Started Guide*, use this guide to learn about basic features common across all OnCommand Insight modules.

OnCommand Insight Assurance User Guide Provides an overview of how to use OnCommand Insight to analyze and validate your storage network environment and to automate monitoring of operations, as well as procedures for making safe changes without disrupting availability.

OnCommand Insight Performance User Guide Provides an overview of how to use OnCommand Insight to reclaim under-utilized resources, manage tiers, identify multipath risks, and troubleshoot ongoing performance bottlenecks.

OnCommand Insight Planning User Guide Provides an overview of how to use OnCommand Insight to enable educated capacity management decisions by managing the end-to-end resource order and allocation process.

Describes how system administrators can communicate storage capacity requirements to SAN managers using OnCommand Insight Connect Applications web access instead of the full OnCommand Insight feature set.

OnCommand Reporting Guide

Describes reports from the centralized reports portal that support viewing critical inventory and capacity-related information from the data warehouse.

Describes the metadata model upon which OnCommand Insight reports are based.

Information in this guide is also available from the OnCommand Insight Data Warehouse portal and in the OnCommand Insight Reporting Connection.

Note: This guide differs from the *OnCommand Report User Guide*, which describes how to use OnCommand Report.

OnCommand Insight Data Warehouse Administration Guide

Describes a data warehouse repository that consolidates multiple OnCommand Insight operational databases in an easy-to-query format.

OnCommand Insight Connect API Reference

Provides an overview of how to use the API interface to enable integration with other applications, such as reporting and monitoring systems. Helps customers and Independent Software Vendors (ISVs) to develop applications using the OnCommand Insight API interface.

Videos for OnCommand Insight

You can access how-to videos that show you how to perform some actions in OnCommand Insight. The videos are located on the OnCommand Insight Customer Community, a forum for customers.

Many individual topics in the OnCommand Insight documentation include links to specific videos.

Note: To see OnCommand Insight videos, use your web browser to access the Customer Community at the following link: https://communities.netapp.com/community/products_and_solutions/storage_management_software/oncommand-insight/oncommand-insight-videos/content

You need a user name and password for some features of this community. To obtain them, click the **Join Now** link.

The following video topics and more are included in the Community:

- How to tell OnCommand Insight about your tiering
- How to identify over utilized disks in OnCommand Insight
- How to view orphaned storage
- How to set thresholds for storage efficiency
- How to identify candidates for tiering
- How to identify candidates for host virtualization
- How to group violations
- How to set business entity annotations

Setting and managing switch thresholds and performance alerts

The administrator can set thresholds that trigger performance alerts and violations.

OnCommand Insight monitors the activity on each switch port. You can change the default thresholds that constitute policies. If these policies are violated, an alert is issued.

Adjusting switch thresholds

You can adjust the default switch thresholds to establish your performance policies. If these thresholds are exceeded, OnCommand Insight generates performance alerts.

Steps

1. From the OnCommand Insight **Open** menu, select **Assurance > Switch Port Performance Alerts**.
2. Right-click on a threshold and select **Configure**.
3. In the Settings window, set the minimum and maximum switch thresholds and the period in minutes.

Switch Thresholds

Set switch performance thresholds based on what the switch is connected to.

Host

Switch

Storage

Generic Device

Tape

Tx Utilization (%)

☒ Min: 2

Period (minutes): 1

☒ Max: 50

Rx Utilization (%)

☒ Min: 2

Period (minutes): 1

☒ Max: 40

Utilization (%)

☒ Min: 2

Period (minutes): 1

☒ Max: 60

CRC (%)

☒ Max: 60

Period (minutes): 1

Other Errors (%)

☐ Max: 10

Period (minutes): 60

Loss of Sync

☐ Max: 1

Period (minutes): 60

Loss of Signal

☐ Max: 1

Period (minutes): 60

Class 3 Discards

☒ Max: 60

Period (minutes): 1

Frame Size Too Long

☒ Max: 60

Period (minutes): 1

Frame Size Too Short

☒ Max: 60

Period (minutes): 1

BB Credit Errors

☐ Max: 1

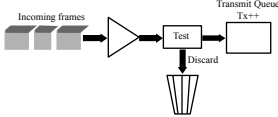
Period (minutes): 60

Apply

Switch threshold types and formulas

OnCommand Insight uses these formulas to determine the switch port performance threshold data.

Threshold	Description	Formula
BB Credit Errors	Fibre Channel uses buffer-to-buffer credits to control transmission flow. The credit value is decremented when a frame is sent and replenished when a response is received. As the available credits for a given port approach zero, the error warns that the port will stop receiving transmissions when zero is reached and will not resume until the BB credits can be replenished.	

Threshold	Description	Formula
Class 3 Discards	The count of Fibre Channel Class 3 data transport discards.	
CRC Rate	<p>CRC Rate is the measure of CRC frame errors as a percentage of the total data traffic.</p> <p>CRC frame errors indicate bit errors somewhere in the data path and point to poor connections, bad cables, or links that are too long.</p>	$\text{crcErrorRate} = \frac{\Delta \text{crcErrors}}{\Delta \text{RxFrames}} \cdot 100$ 
Errors Rate	Total number of errors (Loss of Sync, Loss of Signal, and Framing)	$\text{errorRate}(\%) = \frac{\Delta \text{error}}{\Delta \text{error} + \Delta \text{RxFrames} + \Delta \text{TxFrames}} \cdot 100$
Frame Size Too Long	The count of Fibre Channel data transmission frames that are too long.	
Frame Size Too Short	The count of Fibre Channel data transmission frames that are too short.	
Loss of Signal	If a Loss of Signal error occurs, there is no electrical connection and a physical problem exists.	$\text{lossOfSignal} / \text{sec}(\%) = \frac{\Delta \text{lossOfSignal}}{\Delta \text{sec}} \cdot 100$
Loss of Sync	<p>If a Loss of Sync error occurs, the hardware cannot make sense of the traffic or lock onto it. All of the equipment might not be using the same data rate or the optics or physical connections might be of poor quality.</p> <p>The port must re-sync after each such error, which impacts system performance.</p>	$\text{lossOfSync} / \text{sec}(\%) = \frac{\Delta \text{lossOfSync}}{\Delta \text{sec}} \cdot 100$
Received Utilization	Percentage of available bandwidth used for Rx.	$\text{RxUtilization}(\%) = \frac{\Delta \text{RxBits}}{\Delta \text{Seconds} \cdot \text{ActualSpeedBits}} \cdot 100$
Transmit Utilization	Percentage of available bandwidth used for Tx.	$\text{TxUtilization}(\%) = \frac{\Delta \text{TxBits}}{\Delta \text{Seconds} \cdot \text{ActualSpeedBits}} \cdot 100$

Threshold	Description	Formula
Utilization	Percentage of available bandwidth used for Tx and Rx.	$Utilization(\%) = \frac{\max(\Delta RxBits, \Delta TxBits)}{\Delta Seconds \cdot ActualSpeedBits} 100$

Changing or disabling switch thresholds

You can change switch thresholds at any time. The threshold set at a particular time appears when an alert condition occurs. You can also disable thresholds.

Steps

1. From the OnCommand Insight **Open** menu, select **Assurance > Switch Port Performance Alerts**.
2. In the **Switch Port Performance Alerts** main view, right-click on a threshold and select **Configure**.
The Switch Thresholds settings dialog box opens.
3. Change the values as needed.
4. To disable any thresholds for which you no longer want to receive alerts, clear the **Max** and **Min** check boxes.
5. Click **Apply** to save the changes on one tab and move to another tab.
6. Click **OK** to save all of the threshold changes.

Managing performance alerts

An alert is triggered when the performance metrics for any port exceed the threshold for the time specified. You can review these alerts and identify the switch and port on which the alert condition occurred.

Steps

1. From the OnCommand Insight **Open** menu, select **Assurance > Switch Port Performance Alerts**.
2. In the **Switch Port Performance Alerts** main view, select one or more entries in red.
3. Click the **Alerts** icon.
4. Select one or more entries in the Alerts list.

The screenshot displays two panels from the OnCommand Insight interface. The top panel, titled 'Switch Port Performance Alerts (50) Group: 5', shows a table with columns: Connected To, Threshold, Alerts, Status, Last Time, Minimum, and Max. It lists alerts for a 'Generic Device (10)' and a 'Host (10)', including metrics like Error Rate (%), Loss of Signal, Class 3 Discards, CRC Rate (%), Frame Too Long, Frame Too Short, Received Utilization (%), and Transmitted Utilization (%). The bottom panel, titled 'Alerts (130)', shows a table with columns: Switch, Port, Connected to, Start Time, End Time, Value, Minimum, and Max. It lists alerts for various switches (brocade_gamma, F0lbrod-a) and ports (fc1, fc2, fc3), including metrics like Host exchange_rate, Host ni_exch002, Host ni_orat_bld, Host ni_sap17, Host ni_exch002, Host ni_orat_bld, Host ni_sap17, and Host kic_orat17.

Connected To	Threshold	Alerts	Status	Last Time	Minimum	Max
Generic Device (10)		0			2	
Host (10)		966 8/8 12:37 PM - 8/9 2:17 AM			2	
Host	Error Rate (%)	0				
Host	Loss of Signal	0				
Host	Class 3 Discards	130 8/8/11 12:37 PM			2	
Host	CRC Rate (%)	130 8/8/11 12:37 PM			2	
Host	Frame Too Long	130 8/8/11 12:37 PM			2	
Host	Frame Too Short	130 8/8/11 12:37 PM			2	
Host	Received Utilization (%)	130 8/8/11 2:17 AM			2	
Host	Transmitted Utilization (%)	130 8/8/11 2:17 AM			2	

Switch	Port	Connected to	Start Time	End Time	Value	Minimum	Max
brocade_gamma	fc1	Host exchange_rate, port 10:80:00:00:00:00:00:18	8/8/11 12:00 PM	8/8/11 1:30 PM	1.37	2	
F0lbrod-a	fc1	Host ni_exch002, port 10:80:00:00:00:00:02:52	8/8/11 12:30 PM	8/8/11 1:00 PM	1.50	2	
F0lbrod-a	fc2	Host ni_orat_bld, port 10:80:00:00:00:00:02:54	8/8/11 12:30 PM	8/8/11 1:00 PM	1.00	2	
F0lbrod-a	fc3	Host ni_sap17, port 10:80:00:00:00:00:02:56	8/8/11 12:00 PM	8/8/11 2:00 PM	1.58	2	
F0lbrod-a	fc1	Host ni_exch002, port 10:80:00:00:00:00:02:53	8/8/11 12:00 PM	8/13/11 2:00 AM	0.14	2	
F0lbrod-a	fc2	Host ni_orat_bld, port 10:80:00:00:00:00:02:55	8/8/11 12:00 PM	8/8/11 1:30 PM	1.51	2	
F0lbrod-a	fc3	Host ni_sap17, port 10:80:00:00:00:00:02:57	8/8/11 12:30 PM	8/8/11 12:30 PM	1.52	2	
F0lbrod-a	fc1	Host kic_orat17, port 10:80:00:00:00:00:02:58	8/8/11 12:30 PM	8/8/11 1:00 PM	1.12	2	

- To determine the cause of the alerts, click the icons for any of these views:
 - Switch Port Performance** to evaluate traffic distribution and port utilization.
 - Performance Chart** to identify any issues in performance trends over a specific duration.
 - Changes** to identify any system issues that might have initiated this alert.

Reviewing and confirming alerts

After you identify the cause of an alert or determine that you no longer need to be reminded about that alert, you need to clear alerts from the Switch Port Performance Alerts main view and Alerts detail view.

About this task

An alert is triggered when the performance metrics for any port exceed the threshold for the time specified.

Steps

- In the OnCommand Insight **Open** menu, select **Assurance > Switch Port Performance Alerts**.
- Expand the alert categories and select an alert that you want to examine in more detail.
- Click the **Alerts** icon to see all of the switches and ports associated with the selected alert.
- Select one or more of the switches in the Alerts view.
- Click the **Performance Chart** icon to see a graphic display of the alerts.
- You might also want to display the Changes and Switch Port Performance information for selected alerts.
- To confirm and delete an alert from the display, right-click the line in the Alerts view.

If you do not confirm an alert, it appears every time you display the Switch Port Performance Alerts view.

Examining performance

You can examine your storage environment performance and identify under-utilized and over-utilized resources, maintain tiers, and identify risks before they turn into problems.

OnCommand Insight monitors SAN and NAS devices and does not require any host agents. This performance monitoring does not affect the storage network data paths and never reconfigures devices.

OnCommand Insight helps you to solve or prevent performance and availability problems that are revealed through the collected storage data. In addition, near real-time performance data is available for analysis, using critical array metrics that are kept for seven days.

Note: To see a video demonstrating the OnCommand Insight performance management opportunities, click this link to start the video on the NetApp Community site: [End-to-end performance visibility video](#).

You can use OnCommand Insight to perform these performance management tasks:

- Monitor performance across the environment
- Identify performance trends
- Locate under-utilized volumes
- Balance resources

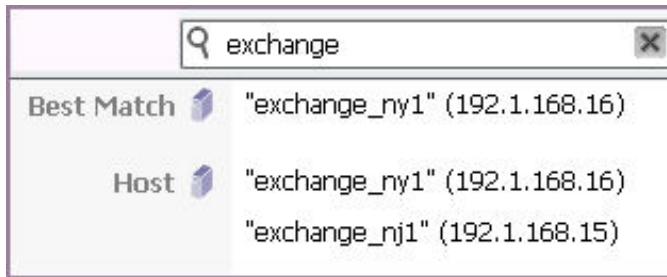
Searching for performance data

You might use the Search facility to display performance data for a portion of your environment that appears to be causing problems. For example, suppose an administrator in the energy business receives a call from a field office reporting that the Sea Current Generation application is running slowly on the server for an energy exchange in New York. You can enter the host name or portion of the name into the Search box to display a summary of the performance information.

Steps

1. In the Search box at the top of the Client window, type at least three characters of the host name.

The search begins immediately. In this example, the administrator enters `exchange` to see a list of all objects with "exchange" in the name.



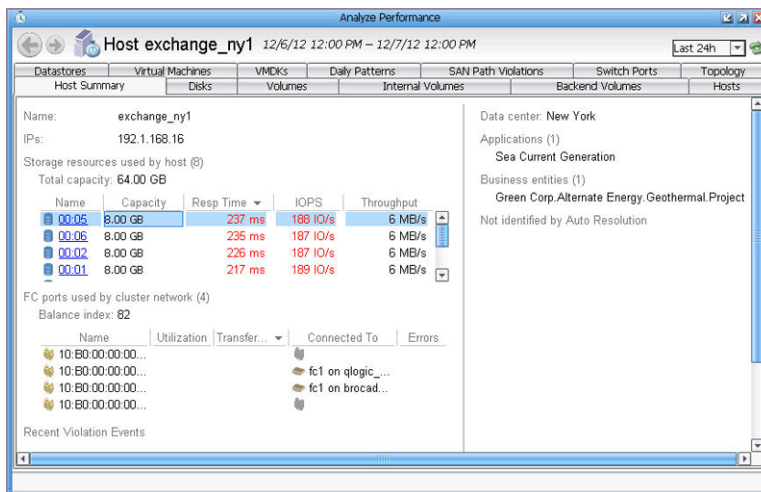
The search results list a New York host (exchange_ny1) and a New Jersey host (exchange_nj1).

2. To examine the performance details for the New York host, double-click the New York host in the search results.

The Host Summary tab of the Analyze dialog box opens to show the description of the host and list problems with the storage resources.

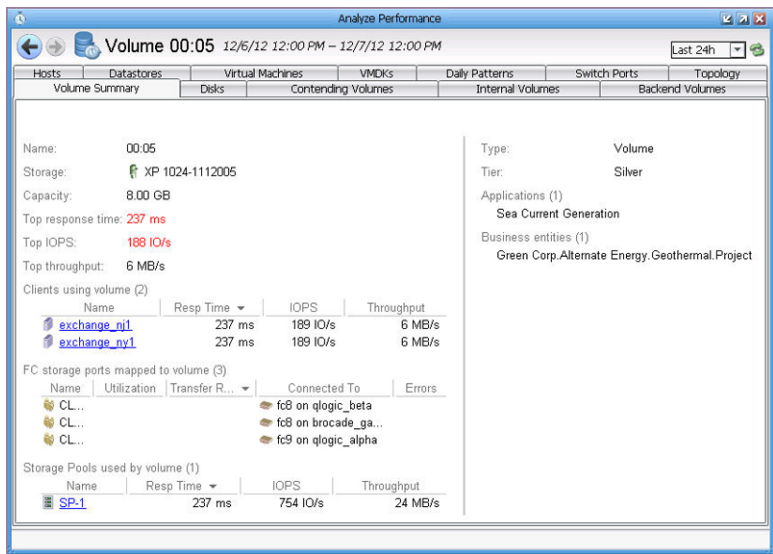
3. In the **Host Summary** tab, check that the correct host name is listed in the description and note what might be wrong.

In this example, the administrator confirms that the selected host is in the New York data center and is running the Sea Current Generation application. In the list of eight Storage Resources, all of the resources show response times in red indicating that they are higher than the performance policy threshold.



4. On the **Host Summary** tab, click the link to any storage resource that you want to examine.

In this example, the administrator clicked the 00:05 resource because it had the highest response time at 237 ms.



5. In this example in the **Volume Summary** tab, the administrator notices that the **Top IOPS** value on this volume is beyond the threshold (shown in red) and that the New Jersey exchange is also a client of this volume. Click the **Contending Volumes** tab to examine the **Volume Performance Chart** for each volume involved in the performance problem.

After you finish

Identify applications that might need to be moved to different volumes or other methods to reduce the performance problem.

Related tasks

[Determining if a VM is affecting host performance](#) on page 41

Related references

[Host Summary tab](#) on page 75


[Volume Summary tab](#) on page 79

Performance trend analysis


OnCommand Insight provides three tools to visualize performance trends: Performance Chart, Distribution Chart, and Daily Performance Pattern.

The **Performance Chart** shows performance trends for one object over time. Use this chart to examine the various performance metrics over the specified period and narrow your search using the detailed information displayed when you position the mouse pointer over points in the chart. To


analyze the specific point in the chart described in the information box, right-click and select **Analyze**. If you select **Show Change Events**, letters display on the chart with links to specific information in the Changes list. In this example, the "A" marks when a data source was added.



The **Distribution Chart** displays trends for several selected objects for visual comparison.



The **Daily Performance Pattern** graph summarizes a metric over the course of a day. The graph illustrates when there are spikes or troughs during a single day and identifies the average hourly values for the analysis period. For switch ports, you can select multiple ports and see the collective impact of traffic on the switch or check individual ports for periods of high or low traffic.



Identifying reasons for performance changes

You can identify possible reasons for performance changes that might be related to configuration changes and use the change markers as research points in a Performance Chart that is linked to the detailed change information.

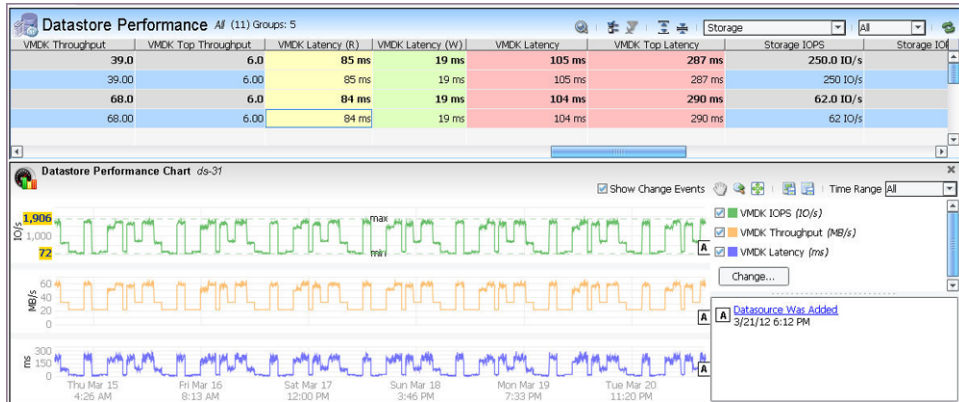
Steps

1. To display the Performance Chart with change markers, open any of these main views:
 - Host Performance
 - Virtual Machine Performance
 - Datastore Performance
 - Switch Port Performance
 - Storage Performance
2. Select a device for which you want to examine performance details.
3. Click the **Performance Chart** icon for the displayed performance view.

Be certain that **Show Change Events** is checked to display change markers if that type of data is available for the view.

Example

This Performance Chart, displayed from Datastore Performance view, includes a change marker with a link to the change details.



- To show the maximum and minimum values for part of the chart, position the mouse pointer over one of the metrics in the Performance Chart legend to the right of the chart.

Example

In this example, the maximum and minimum values for the VMDK IOPS are shown and the “A” change marker indicates when a data source was added.

- Click on the link beside the change marker letter to display the **Changes** detail view listing the configuration changes that occurred at that time.

Related tasks

[Analyzing virtual storage performance](#) on page 50

Related references

[Performance Chart](#) on page 101

Monitoring system changes that impact performance

You can identify system changes that might impact performance. Multiple changes generated from a single event are listed with the same date and time. To reduce the number of messages displayed for any one event in the **Changes** view, clear the **Show transient violations** check box in the **Tools > Settings** for the **General Preferences**.

About this task

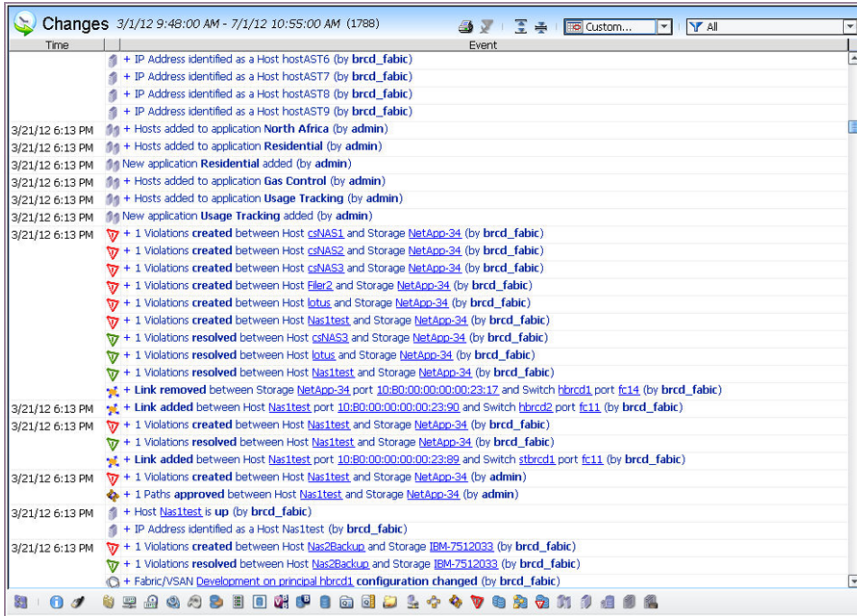
You can view the following types of physical and logical changes to SAN devices:

- Equipment additions and removals
- Zoning and volume masking changes
- Cabling reconfiguration

- Outages
- Policy changes

Steps

1. From the OnCommand Insight **Open** menu, select **Assurance > Changes**.



2. Scroll through the list of changes to identify events that might be causing problems and click the links in the descriptions to see more details.
3. If you are concerned about performance relating to the switch ports, you can open **Assurance > Switch Port Performance Alerts** and select one or more items listed in red.
4. Click the **Alerts** icon to display information about the selected switches, ports, and their connections.
5. If the switch thresholds need to be adjusted, right-click an item in the **Switch Port Performance Alerts** view and select **Configure** to change the threshold settings.

Related tasks

[Troubleshooting slow client computers on your network](#) on page 43

[Troubleshooting poor application performance](#) on page 44

Monitoring Fibre Channel switch performance

You can examine the performance of the Fibre Channel switches in your environment and identify risks before they turn into problems.

Steps

1. To analyze Fibre Channel switch performance, select the OnCommand Insight **Open** menu, select **Performance > Switch Port Performance**.
2. You can use "No Grouping" or select a grouping for the data from these menu options:
 - Fabric then Switch
 - Connected Device Type and then Name
 - Connected Device
 - BU then Application then Connected Device
 - Application then Connected Device
 - Switch then Connected Device
 - Switch then Blade
 - Storage then Controller

Example

This Switch Port Performance view shows the traffic moving through switch ports. In this example, the data is grouped by Connected Device Type and then Name.

Switch	Switch Port	Has Errors	Value	Rx Value	Tx Value	Distribution	Fabric/VSAN	VSAN	Speed
Host (162)		0	4,674,115.73 GB	3,972,92...	701,187...	57.6%			
Storage (162)		0	3,120,159.09 GB	2,652,08...	468,076...	38.4%			
Switch (26)		0	271,921.65 GB	231,127...	40,793.9...	3.4%			
Switch 72 (2) Switch 72		0	64,023.32 GB	54,418.4...	9,604.87 ...	23.5%	20:00:00:...		4G
Switch 72	fc0		58,151.69 GB	49,428.24...	8,723.45 GB	90.8%	20:00:00:0...		4G
Switch 72	fc2		5,871.63 GB	4,990.22 GB	881.41 GB	9.2%	20:00:00:0...		4G
Switch 72 (2) Switch 73		0	63,750.08 GB	54,186.2...	9,563.85 ...	23.4%	20:00:00:...		4G
Switch 73	fc0		57,955.78 GB	49,261.75...	8,694.03 GB	90.9%	20:00:00:...		4G
Switch 73	fc2		5,794.30 GB	4,924.47 GB	869.82 GB	9.1%	20:00:00:0...		4G
hdis300 (2) scis300		0	32,213.36 GB	27,380.9...	4,832.41 ...	11.8%			4G
scis300	fc2/14		16,193.21 GB	13,764.03...	2,429.17 GB	50.3%			4G
scis300	fc2/15		16,020.15 GB	13,616.92...	2,403.24 GB	49.7%			4G
hdis200 (2) scis200		0	24,496.01 GB	20,821.2...	3,674.78 ...	9.0%			4G
scis300 (2) scis300		0	17,410.65 GB	14,798.6...	2,612.01 ...	6.4%			4G
Switch 78 (1) Switch 79	fc0	0	16,054.57 GB	13,646.1...	2,408.38 ...	5.9%	20:00:00:...		4G
hbrcd1 (1) stbrcd1	fc15	0	10,183.44 GB	8,655.84 ...	1,527.60 ...	3.7%	20:00:00:...		2G
hbrcd2 (1) stbrcd1	fc0	0	10,093.92 GB	8,579.75 ...	1,514.17 ...	3.7%	20:00:00:...		2G
F3brcd-b (1) F3brcd-a	fc1	0	9,970.87 GB	8,475.14 ...	1,495.73 ...	3.7%	20:00:00:...		2G

3. In the Switch Port Performance view, select a device and right-click to select any of these options:

- **Show Performance Histogram** displays the relative amount of time the selected ports spent in the given utilization levels.
 - **Show Daily Traffic Pattern** identifies the flow of traffic by hour, day, week or by a custom calendar setting.
 - Analyze enables the correlation of data and shows trend and pattern changes.
 - Analyze Contention enables you to analyze the reasons for poor performance and identify where that performance has degraded.
4. To view a chart of the performance for an item in the Switch Port Performance view, select a device or group of devices and click the **Port Performance Distribution** icon.

Example

In this example, the selected devices are shown as individual colors in the distribution chart. The legend on the right lists the colors associated with the devices.

5. Position the mouse pointer on points of interest in the Port Performance Distribution to display additional time and device details.
6. To see more details, click the **Performance Chart** and **VM Distribution** icons.

Monitoring data store performance

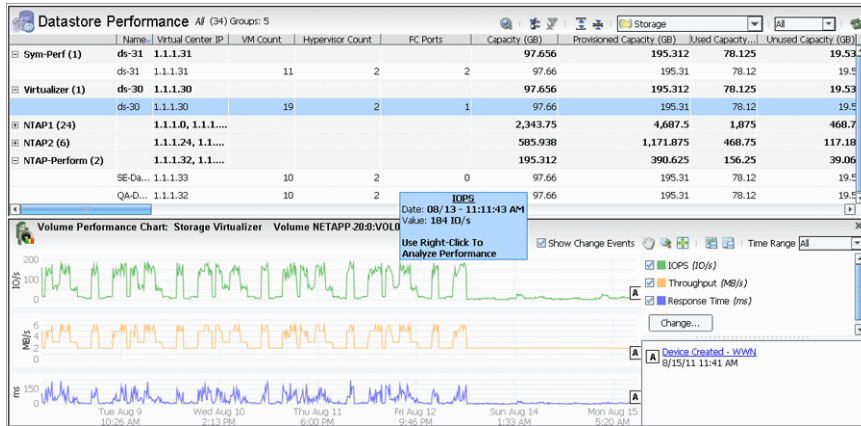
You can monitor data store performance characteristics (such as IOPS, throughput, latency CPU and memory utilization) and compare storage performance with data store performance in the performance charts. This information helps to identify busy data stores on the V-Cluster and balance data stores after reviewing the data store load and storage load.

Steps

1. To analyze data store performance, select the OnCommand Insight **Open** menu, select **Performance > Datastore Performance**.
2. Group the data display and select a time period. Examine the data for any anomalies.
3. Select any data stores of concern and click the **Storage Resource Performance Chart** icon.
4. To display detailed information for any point in the graph, you can position the mouse pointer over a point in the chart and click or click the **Pan** (hand) icon to display details quickly for each point as you move the mouse pointer over the graph.

Example

In this example, the details for an IOPS peak are displayed.



5. To analyze the performance of a selected point in a chart, right-click, and select **Analyze**.
6. In the **Datastore Summary** tab, click the linked **Name** of the volume to display more detailed information.
7. Click the **Virtual Center IP** link to go to Virtual Center.
8. To examine information from the related VMDKs, click the **VMDKS** tab.

Related references

Data Store Summary tab on page 72

Monitoring storage performance

After you set performance thresholds in the OnCommand Insight Settings, you can use the Storage Performance main view to observe disk utilization, IOPS, and volume response time to identify potential problems before a crisis.

Before you begin

The Array Performance thresholds must be established in the Global Policies section of the **Tools > Settings**.

Steps

1. Open the **Storage Performance** main view and group the information according to the focus of your research.

In this example, the data is grouped by Data Center.

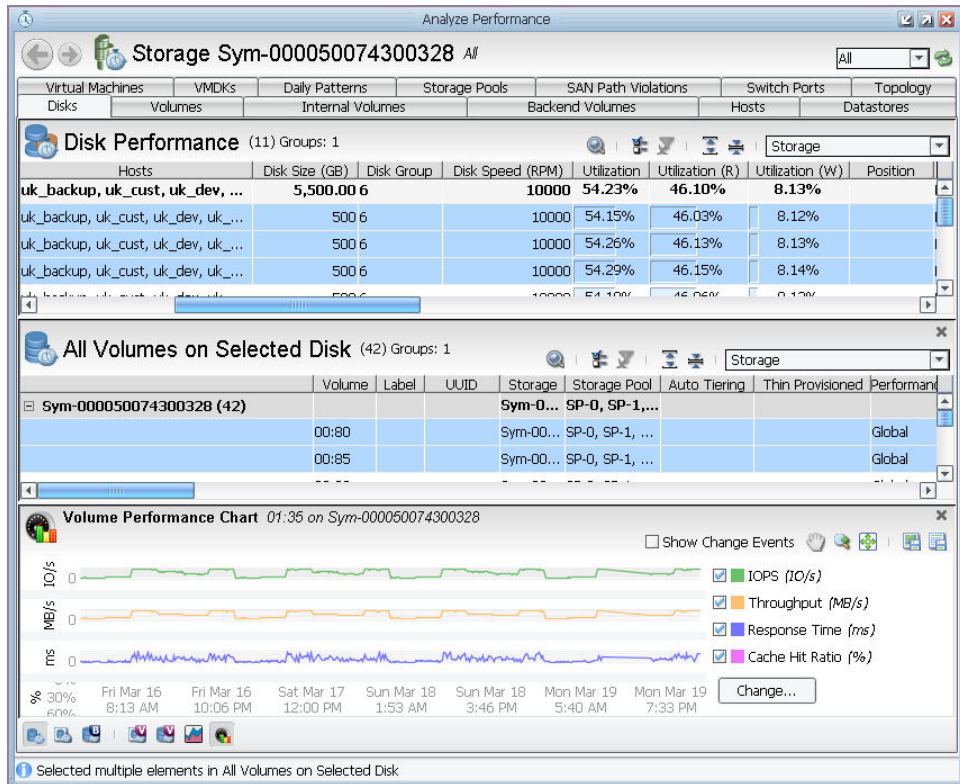
2. Expand the groups to show the contents of each and note the virtualization types, capacities, and response times.

Storage Performance All (19) Groups: 4										
	Name	Technology	Capacity (GB)	Raw Capacity (GB)	Vendor	Family	Model	Volume Response Time	Volume Response Time (R)	Volume Re
Tokyo (5)			22,318.00	34,688.00 EMC, NetApp						
	Sym-000050074300318		3,750.00	5,000.00 EMC	Sym...	Sym...		83 ms	81 ms	
	NTAP-Perform	FC,SCSI,CI...	5,120.00	10,240.00 NetApp	FAS6...	FAS6...				
	Sym-000050074300327		3,000.00	6,000.00 EMC	Sym...	Sym...		93 ms	92 ms	
	Sym-000050074300328		3,000.00	6,000.00 EMC	Sym...	Sym...		92 ms	90 ms	
	CX600_29_CK00351029329	FC	7,448.00	7,448.00 EMC	Clarion	CX600		0 ms	0 ms	
Zurich (6)			126,712.54	185,956.00 EMC, NetApp						
	APM000910025			EMC	Celer...	Celer...				
	Symmetrix-FAST		113,476.54	162,868.00 EMC	Sym...	Sym...		84 ms	82 ms	
	V-Series-Zurich2	FC,SCSI,CI...	64.00	NetApp	V-Series	V-Series		68 ms	67 ms	
	V-Series-Zurich1	FC,SCSI,CI...	64.00	NetApp	V-Series	V-Series		70 ms	69 ms	
	Sym-000050074300324		3,000.00	5,000.00 EMC	Sym...	Sym...		78 ms	76 ms	
	CX600_26_CK00351029326	FC	10,108.00	18,088.00 EMC	Clarion	CX600		0 ms	0 ms	
Sydney (4)			17,076.00	30,864.00 EMC, NetApp				83 ms	81 ms	
	CX600_31_CK00351029331	FC	4,788.00	6,394.00 EMC	Clarion	CX600		0 ms	0 ms	
	Sym-Perf		2,000.00	4,000.00 EMC	Sym...	Sym...		80 ms	79 ms	
	NTAP-Perf	FC,SCSI,CI...	10,240.00	20,480.00 NetApp	FAS3...	FAS3...		80 ms	78 ms	
	Virtualizer	FC,SCSI,CI...	48.00	NetApp	V-Series	V-Series		83 ms	81 ms	
New York (4)			21,118.00	30,108.00 EMC, HDS, HP				86 ms	85 ms	
	CX600_32_CK00351029332	FC	6,118.00	10,108.00 EMC	Clarion	CX600		0 ms	0 ms	
	USP-2222112	FC	6,250.00	7,500.00 HDS	USP	USP		0 ms	0 ms	
	USP-2222111	FC	6,250.00	7,500.00 HDS	USP	USP		0 ms	0 ms	
	XP 1024-1112005	FC	2,500.00	5,000.00 HP	XP	XP 10...		86 ms	85 ms	

3. Scroll across the data to examine storage of interest and its Volume Cache Hit ratios and Disk Utilization percentages, as shown below.

Storage Performance All (19) Groups: 4										
	Throughput (W)	Top Volume Throughput	Volume Cache Hit Ratio	Volume Cache Hit Ratio (R)	Volume Cache Hit Ratio (W)	Volume Partial R/W	Write Pending	Volume Accessed	Internal Volume Be	
	15 MB/s	6 MB/s	45%	45%	44%	20%	0	Yes		
	243 MB/s	92 MB/s	44%	44%	44%	20%	0	Yes		
	243 MB/s	92 MB/s	45%	44%	45%	20%	0	Yes		
	42 MB/s	2 MB/s	0%	0%	0%	0%	0	Yes		
	2,378 MB/s	1,390 MB/s	44%	44%	44%	20%	0	Yes		
	1 MB/s	6 MB/s				20%	0	Yes		
	1 MB/s	6 MB/s				20%	0	Yes		
	20 MB/s	6 MB/s	44%	44%	44%	20%	0	Yes		
	10 MB/s	2 MB/s	0%	0%	0%	0%	0	Yes		
	47 MB/s	87 MB/s				15%	0	Yes		
	9 MB/s	2 MB/s	0%	0%	0%	0%	0	Yes		
	25 MB/s	87 MB/s	44%	44%	45%	20%	0	Yes		
	7 MB/s	6 MB/s				20%	0	Yes		
	6 MB/s	6 MB/s				20%	0	Yes		
	82 MB/s	6 MB/s	11%	11%	11%	5%	0	Yes		
	8 MB/s	2 MB/s	0%	0%	0%	0%	0	Yes		
	34 MB/s	2 MB/s	0%	0%	0%	0%	0	Yes		
	32 MB/s	2 MB/s	0%	0%	0%	0%	0	Yes		
	8 MB/s	6 MB/s	45%	45%	45%	20%	0	Yes		

4. Select one or more storage arrays, right-click, and then select **Analyze**.
5. Use the tabs to see data for different devices associated with the storage and click icons to display more detailed information for your selections, as shown here.



- If you conclude that the performance policy thresholds should not be used for individual arrays, you can use the **Modify Policy** wizard to customize the global policy for those specific storage arrays and the storage pools.

Customizing global general policies

You might want to modify global general policies, not the global SAN path policies, for special characteristics in one area of your environment.

About this task

Using the Modify Policy wizard, you can customize the global general policies, create exceptions to these policies, and remove exceptions to the global policies for the selected items. You can also disable the evaluation of the parent global policy at the selected level.

Steps

- From OnCommand Insight **Open** menu, select **Assurance** and any of these views:
 - Violations Browser
 - Disk Utilization Violations

- Port Balance Violations
- Policies

Or you might open any of these **Performance** views:

- Host Performance
- Datastore Performance
- Storage Performance

Or you might open any of these **Inventory** views:

- Hosts
- Datastores
- Storage Arrays
- Tapes

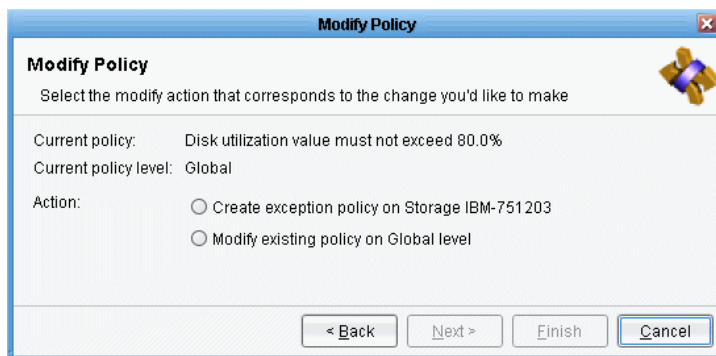
2. Select one or more items in the selected view to examine and possibly modify the global general policy to customize it for one area of your environment.
3. Right-click and select **Modify Policy**.
4. If more than one policy type applies to the selected items, these policy types are listed on the first page of the Modify Policy wizard. Select the policy type you want to examine and possibly change. Click **Next**. Otherwise, the second page is displayed.

At any point, you may return to previous pages in the wizard to review or make additional changes.

5. Select the type of policy modification you want.

Example

In this example, there are two mutually exclusive options that might be selected for the global disk utilization policy type.



If you selected more than one item in the view, you might be modifying different settings of the same policy, and the wizard shows that condition by listing the **Current policy** as "Mixed."

6. Select a policy.
7. Click **Next** to save the selected modification type.
8. The third page of the wizard allows you to enter settings for the policies. Make the setting selections and click **Next**.
9. The fourth and final wizard page lists the previous policy and the changes you made so that you can review and confirm those changes.

If there were multiple policy types listed on the first page and you want to make changes to a different policy type, select **Modify another policy on the same selections** on the confirmation page.

10. Click **Finish** to save your changes

Note: All changes to the policies are recorded in the Audit log.

General policy types and their hierarchies

The general policies that evaluate the elements in your environment are governed by the policy hierarchy. All of the global general policy types, listed below, have default policies that govern your environment. You can customize the default global policies and add policy exceptions on lower-level (not global level) components in the hierarchy.

Blocked Generic Devices	Alerts the administrator when unidentified hosts (generic devices) cannot contact any volumes or shares.
Blocked Hosts	Alerts the administrator when a host cannot contact any volumes or shares. This policy can only be set at the global level.
Datastore Latency	Sets the virtual machine performance threshold as the maximum latency of the data stores in your environment. You can set the notification for the violation to be generated when the latency first exceeds the threshold (peak) or when the average of an hour of samples exceeds the threshold. You can create a policy exception for a data store.
Disk Utilization	Sets the threshold for the hourly average of the disk utilization percentage. You can create a policy exception for a data store.
Host Port Balance	Sets the threshold for the traffic load across a device's Fibre Channel (FC) ports to be evenly distributed. You can create a policy exception for the storage level.
Internal Volume Performance: IOPS	Sets the threshold for internal volume maximum IOPS and allows you to send notifications when the IOPS reach the threshold. You can create a policy exception for the storage array and then for the internal volume.
Internal Volume Performance: Response Time	Sets the threshold for the maximum internal volume response time. You can create a policy exception for the storage array and then for internal volumes.

Internal Volume Performance: Throughput	Sets the threshold for internal volume maximum throughput and allows you to send notifications when the throughput reaches the threshold.
Storage Pool Used Capacity	Sets the percentage threshold for the used capacity of the storage pool. You can create a policy exception for storage array and then for storage pools.
Storage Pool Capacity Assurance	Defines the thin-provisioned storage pools (aggregate) where the sum of unused capacity exceeds current unused space of the storage pool. A violation is generated when a storage pool does not have enough unused capacity to accommodate remaining reserved unused space. By Default, 3 Volumes need to maintain enough capacity to accommodate unused space. Therefore, the number in the global policy is the number of volumes required to maintain enough capacity to accommodate unused space.
Storage Pool Commit Ratio	Sets the percentage threshold for the commit ratio. You can create a policy exception for storage and then for storage pools.
Storage Port Balance	Sets the threshold for the traffic on storage FC ports to be evenly distributed. You can create a policy exception for storage array.
Storage Pool Fan-Out	Sets the number of masked hosts (generic devices) for storage ports. This policy can only be set at the global level.
Tape Port Balance	Sets the threshold for the traffic on a tape's FC ports to be evenly distributed. You can create a policy exception for tape.
Volume Performance: IOPS	Sets the threshold for volume maximum IOPS and allows you to send notifications when the IOPS reach the threshold. You can create a policy exception for the storage array and then for the volume.
Volume Performance: Response Time	Sets the threshold for the hourly average or the peak of the volume response time. You can create a policy exception for the storage and then for the volume.
Volume Performance: Throughput	Sets the threshold for volume maximum throughput and allows you to send notifications when the throughput reaches that threshold.

Identifying performance trends

You can examine the OnCommand Insight performance data to see correlations between performance and errors as well as anomalies and sudden changes in system activity to identify the trends.


Steps

1. To identify performance trends, open any of these performance views:

- Host Performance
 - Virtual Machine Performance
 - Datastore Performance
 - Switch Port Performance
 - Storage Performance
2. Select an element for which you want to see performance details.
 3. Click the **Performance Chart** icon.
The Performance Chart legend associates the types of metrics with their colors. You can select one or more metrics.
 4. To change the period, select the time from the **Time Range** menu in the upper right corner.
 5. Some of the metrics legends include a Change button. Click that button to select different metrics to display on the Performance Chart.
 6. You might also want to position the mouse pointer over chart points to see more information for those points in the chart and click items in the legend to show the minimum and maximum range in the selected value on the chart.

Example

This example shows the Performance Chart displayed for selected Storage Performance items with the minimum and maximum Response Time highlighted from the legend.



Related references

[Storage Performance view](#) on page 106

[Performance Chart](#) on page 101

Exporting performance data as an image or file

You can export any of the Performance or Distribution charts as images and to comma-separated values (csv) files for import to Excel and other tools.

Steps

1. Display performance data in a Performance Chart or Performance Distribution chart for one of the performance views.
2. In the upper right corner of the view, click one of these icons:
 - **Export as Image to File**
 - **Export Raw Data to CSV File**
3. In the Windows dialog box, select a directory and enter the name of the file you are creating.
4. Depending on the type of file you are creating, click **Save** for the image file or **Export** for the CSV file.

If you exported the data as a CSV file, the format is Date, Element name, and Metric Value.

Related references


- [Performance Chart](#) on page 101
- [Performance Distribution view](#) on page 103

Grouping data in views

You can classify and group the data in a view table so that you can interpret the data in groups that you specify. For example, you can group switches data by fabric.

Steps

1. In the right of any view, click the Grouping drop-down list.



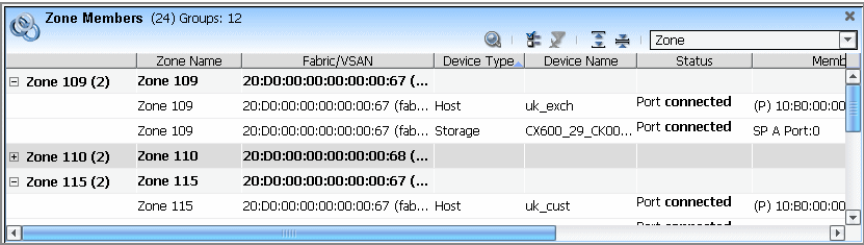
2. Select a data grouping.
3. To remove the grouped display from the table, select "No Grouping" from the Grouping drop-down list.

Note: You can also group data according to annotation groups. Using annotation groups, you can assign custom business-specific categories to resources and then group the resources by those categories. For example, on the Hosts view, you can group the data by names of your data centers so that you can see all resources for each data center.

Example of grouped data

The following example shows the Zone Members detail view with the Zone grouping.

If you do not group data, the data appears on one row per resource. On the other hand, if you do group the data, OnCommand Insight adds a column to the left of the table with grouping information. The number in parentheses after each grouping level represents the number of items reported in that grouped row.



	Zone Name	Fabric/VSAN	Device Type	Device Name	Status	Member ID
[-] Zone 109 (2)	Zone 109	20:00:00:00:00:00:67 (...)				
	Zone 109	20:00:00:00:00:00:67 (fab...	Host	uk_exch	Port connected	(P) 10:80:00:00
	Zone 109	20:00:00:00:00:00:67 (fab...	Storage	CX600_29_Ck00...	Port connected	SP A Port:0
[+] Zone 110 (2)	Zone 110	20:00:00:00:00:00:68 (...)				
[-] Zone 115 (2)	Zone 115	20:00:00:00:00:00:67 (...)				
	Zone 115	20:00:00:00:00:00:67 (fab...	Host	uk_cust	Port connected	(P) 10:80:00:00

Monitoring virtual machine performance

Virtual machine (VM) administrators need to examine contention with VM resources and look for locations for new VMS.

VM administrators can use OnCommand Insight to research the following:

- Determine where to add a new VM.
- Determining if there is any VM contention related to I/O or memory utilization.
- Select an active path for a VM datastore based on low port utilization.
- Add a host to a virtual cluster to share an existing datastore.
- Determine whether a VM is responsible for a host performance issue.

Viewing virtual machine performance and utilization

Storage administrators often need to look for inactive virtual machines or virtual machines that are using high tier storage that could be using a lower tier storage. Using OnCommand Insight, you can view virtual machine performance and utilization. When all virtual disks are in NFS/iSCSI backend data stores, you can also see traffic going to individual virtual disks attached to a virtual machine.

About this task

The Virtual Machine Performance main view does not display information about virtual machines and LUNs that do not map to a volume.

The colors that display on virtual machine performance views reflect the settings on the VM Hosts Thresholds view in **Tools > Settings**.

Steps


1. From the OnCommand Insight **Open** menu, select **Performance > Virtual Machine Performance**.

Name	DNS Name	Host Names	Datastore	V-Cluster	CPU Utilization	Memory Utilization	Disk Latency (R)	Disk Latency (W)	Disk Latency (Top)	Disk IOPS (R)	Disk IOPS (W)
vm-60	vm-60.ona...	esx1	ds-30	Cluster 1	85.4%	96.8%	83 ms	19 ms	102 ms	277 ms	71 IOPS
vm-61	vm-61.ona...	esx3	ds-31	Cluster 2	84.2%	87.4%	87 ms	19 ms	107 ms	284 ms	71 IOPS
vm-62	vm-62.ona...	esx1	ds-30	Cluster 1	8.7%	23.4%	0 ms	0 ms	0 ms	1 ms	0 IOPS
vm-63	vm-63.ona...	esx2	ds-30	Cluster 1	8.9%	23.5%	0 ms	0 ms	0 ms	1 ms	0 IOPS
vm-64	vm-64.ona...	esx3	ds-31	Cluster 2	8.6%	23.3%	0 ms	0 ms	0 ms	1 ms	0 IOPS
vm-65	vm-65.ona...	esx4	ds-31	Cluster 2	8.7%	23.7%	0 ms	0 ms	0 ms	1 ms	0 IOPS
vm-66	vm-66.ona...	esx1	ds-30	Cluster 1	8.7%	23.7%	0 ms	0 ms	0 ms	1 ms	0 IOPS


2. In the **Virtual Machine Performance** main view, optionally, select one of the following options to group the data in the view:

- Host

- Datastore
 - V-Cluster then Host
 - V-Cluster then Datastore
3. To determine whether there are anomalies or sudden changes in the virtual machines over time, click on the **Performance Chart** icon in the status bar.



4. To see traffic distribution on the virtual machine, click the **Performance Distribution** view icon.



Related references

[Performance Distribution view](#) on page 103

Setting virtual machine thresholds

Set the range of Hypervisor (HV) performance threshold values that show up in green, yellow, or red on the Host Performance main view, which displays the HV utilization of virtual machine (VM) hosts (ESX servers) over a selected period.

About this task

For each threshold, you set high and low values. After you set the threshold, severity is represented on the Host Performance main view using the following colors:

- Green: Values below the threshold's low value.

- Yellow: Values below the threshold's high value, but above the threshold's low value.
- Red: Other values, such as those above the threshold's high value.

The settings on this view affect the colors on the following views: Datastore Performance main view, Virtual Machine Performance main view, VMDK view, and Host Performance main view.

Steps

1. From the OnCommand Insight Client menu, select **Tools > Settings**.
2. In the left tree of the **Settings**, click **Thresholds > VM Hosts Thresholds** option.
3. For each threshold set low and high limits:
 - Latency: Average amount of time for a read/write operation from or to the virtual disk before an alert is issued. The default values are Low 20 and High 100.
 - CPU Utilization: Threshold for the host's CPU MHz being used. The default values are Low 45 and High 75.
 - Memory Utilization: Threshold for the memory used by the host. The default values are Low 35 and High 65.
4. To save settings on this one view and continue with other settings, click **Apply**.
5. Click **OK**.

Determining where to add a data store to achieve load balancing

The virtualization administrator (VA) can add a new data store to a single host or a cluster and achieve load balancing.

About this task

The major steps are:

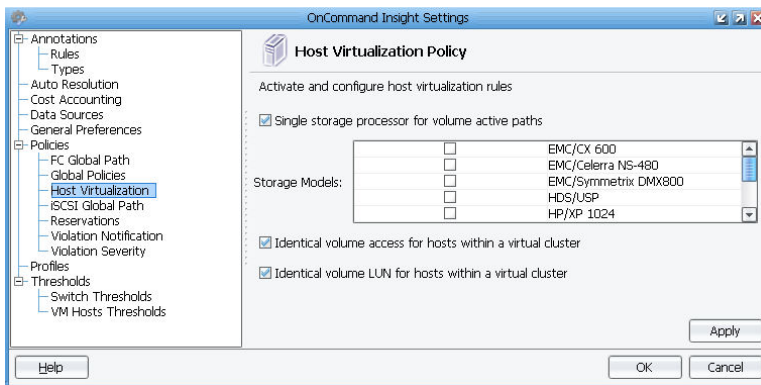
- The VA asks the storage administrator (SA) to allocate the capacity and enable access from the hosts.
- The SA allocates the capacity and configures the storage array and SAN switches to enable multiple redundant paths.
- The VA creates a data store on the allocated volume.
- The VA chooses the preferred paths to achieve load balancing while avoiding path thrashing.

Steps

1. Allocate some capacity and configure access for a group of hosts.
2. To identify existing volume allocations for the hosts, select **Inventory > Hosts**.

This information supports your decision for volume allocation. You might choose storage arrays that are already zoned, for example.

3. Allocate a new volume.
4. To choose the storage port and storage processor that will allow for a balanced load, select **Performance > Switch Port Performance**.
5. Configure the storage array and fabric switches to enable all hosts to access the volume through multiple paths.
6. To verify that all hosts have access to the volume, select **Inventory > Hosts**.
7. Create a data store on the new volume and configure the preferred path from the hosts.
8. Select **Inventory > Virtual Machines**.
9. From the **Virtual Machines** main view, click the Virtual Machines Disks icon in the status bar to show the **VMDKs** detail view.
10. Verify that the data store was added to all hosts in the cluster.
11. To set a global policy, do the following:
 - a) Select **Tools > Settings**.
 - b) In the Policies list, select **Host Virtualization**.



- c) Set the Host Virtualization global policy that will continuously verify that the hosts have sufficient redundant access to the volume and that the conditions for path thrashing do not occur.

Path thrashing might be caused by conflicting selections of preferred paths.

Determining virtual machine contention

If you find that a VM is performing poorly, you can determine whether there is contention for I/O or CPU utilization. You might find that the VM is performing slowly, and you can see whether this is a


standard trend over time or just a spike. You can also look at the data stores and determine which of those are busiest, giving you data to help with future migrations.

Steps

1. Select **Performance > Virtual Machine Performance**.

Name	DNS Name	Host Names	Datastore	V-Cluster	CPU Utilization	Memory Utilization	Disk Latency (R)	Disk Latency (W)	Disk Latency	Top Disk Latency	Disk IOPS (R)	Disk IOPS (W)
vm-60	vm-60.ona...	esx1	ds-30	Cluster 1	85.4%	86.6%	83 ms	19 ms	102 ms	277 ms	71 IOPS	1 IOPS
vm-61	vm-61.ona...	esx3	ds-31	Cluster 2	84.2%	87.4%	87 ms	19 ms	107 ms	284 ms	71 IOPS	1 IOPS
vm-62	vm-62.ona...	esx1	ds-30	Cluster 1	8.7%	23.4%	0 ms	0 ms	0 ms	1 ms	0 IOPS	0 IOPS
vm-63	vm-63.ona...	esx2	ds-30	Cluster 1	8.9%	23.5%	0 ms	0 ms	0 ms	1 ms	0 IOPS	0 IOPS
vm-64	vm-64.ona...	esx3	ds-31	Cluster 2	8.6%	23.3%	0 ms	0 ms	0 ms	1 ms	0 IOPS	0 IOPS
vm-65	vm-65.ona...	esx4	ds-31	Cluster 2	8.7%	23.7%	0 ms	0 ms	0 ms	1 ms	0 IOPS	0 IOPS
vm-66	vm-66.ona...	esx1	ds-30	Cluster 1	8.7%	23.7%	0 ms	0 ms	0 ms	1 ms	0 IOPS	0 IOPS

2. To see which virtual machines are using the most resources, click the **VM Performance Distribution** view icon in the status bar.



3. To see whether this is a standard trend over time or an unusual spike, click the **Performance Chart** icon in the status bar.



4. To analyze performance, in the **Virtual Machine Performance** main view, right-click on a VM and select **Analyze Contention**.
5. In the Analyze Contention view, look for hosts that might be in contention for performance on the disk, a factor that can cause poor performance.

Related references

[Analyze Contention dialog box](#) on page 68

[Performance Chart](#) on page 101

[Performance Distribution view](#) on page 103

Selecting an active path for a virtual machine data store

Use OnCommand Insight to determine which active path to set for a virtual machine (data store). You should look for a low utilization storage port.

About this task

Look for low storage port utilization on the storage array on which the virtual machine was allocated (for a new data store) or for the storage processor that is already being used for access by other ESX servers (for an existing data store and a new ESX server).

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Switch Port Performance**.

Switch	Switch Port	VAS	Distribution	Fabric	VSAN	Port Speed	Switch Blade	Port Controller	Transfer Rate	Connected Device Type	Connected Device Name
Switch (26)		253,060,792 MB	5.0%						1,170.86 M...	Switch	
Storage (125)		2,146,240,113 MB	42.4%						5,974.98 M...	Storage	
San Antonio (2)		616,263 MB	0.0%			20		41	3.49 MB/s	Storage	San Antonio
Switch 36	fc1	12,254 MB	2.0%	20:00:00...		20		41	0.07 MB/s	Storage	San Antonio
Switch 37	fc1	604,009 MB	96.0%	20:00:00...		20		41	3.42 MB/s	Storage	San Antonio
Phoenix (2)		618,446 MB	0.0%			20		36	3.51 MB/s	Storage	Phoenix
Switch 34	fc1	606,047 MB	96.0%	20:00:00...		20		36	3.44 MB/s	Storage	Phoenix
Switch 35	fc1	12,400 MB	2.0%	20:00:00...		20		36	0.07 MB/s	Storage	Phoenix
Philadelphia (2)		1,426,061 MB	0.1%			26		31	8.08 MB/s	Storage	Philadelphia
V-Series-Zurich...		2,698,473 MB	0.1%			26			15.30 MB/s	Storage	V-Series-
Dallas (2)		5,862,106 MB	0.3%			26		16	33.23 MB/s	Storage	Dallas
Jacksonville (2)		5,920,484 MB	0.3%			26		26	33.56 MB/s	Storage	Jacksonville
Sym-00005007...		6,356,391 MB	0.3%			26		99	10.57 MB/s	Storage	Sym-000
NTAP2 (2)		7,321,201 MB	0.3%			26		51	41.50 MB/s	Storage	NTAP2
V-Series-Zurich...		8,138,647 MB	0.4%			26		82	46.14 MB/s	Storage	V-Series-
Detroit (2)		8,232,411 MB	0.4%			26		21	46.67 MB/s	Storage	Detroit
Sym-00005007...		14,095,236 MB	0.7%			26			79.90 MB/s	Storage	Sym-000

2. From the **Switch Port Performance** main view, select "Last week" as the timescale.
3. Filter for storage arrays.
4. Select the storage device with the data store volume or volumes. Select multiple devices.
5. If you need a longer historical view, in the **Switch Port Performance** main view, adjust the time.
6. Click the **Port Performance Distribution** icon in the status bar.
7. Identify storage ports that have the least utilization. Check for connectivity to the same fabric if the host is already connected.

Adding a host to a virtual cluster to share a data store

The virtualization administrator (VA) can add a new host to a group of hosts in a virtual cluster and share an existing data store.

About this task

The major steps for adding a host to a virtual cluster to share a data store are:

- The VA asks the storage administrator (SA) to grant the host access to the SAN volumes on a data store.
- The SA configures the storage array and SAN switches to enable multiple redundant paths.
- The VA adds the data store to the host and chooses the preferred paths to achieve load balancing and avoid path thrashing.

Steps

1. Ask the system administrator to grant you host access rights to the volumes on a data store.
2. To identify which storage ports and processors are being used to access this data store by existing hosts, select **Inventory > Hosts**
3. To identify the storage port and storage processor that will allow for a balanced load and choose the correct mapping to avoid path thrashing, select **Performance > Switch Port Performance**.

	Switch	Switch Port	Has Errors	Value	Rx Value	Tx Value	Distribution	Fabric/VSAN	VSAN	Speed
Host (162)			0	4,674,115.73 GB	3,972,92...	701,187...	57.6%			
Storage (162)			0	3,120,159.09 GB	2,652,08...	468,076...	38.4%			
Switch (26)			0	271,921.65 GB	231,127...	40,793.9...	3.4%			
Switch 73 (2)	Switch 72	fc0	0	64,023.32 GB	54,418.4...	9,604.87 ...	23.5%	20:00:00:...		4G
	Switch 72	fc2		5,871.63 GB	4,990.22 GB	881.41 GB	9.2%	20:00:00:0...		4G
Switch 72 (2)	Switch 73		0	63,750.08 GB	54,186.2...	9,563.85 ...	23.4%	20:00:00:...		4G
	Switch 73	fc0		57,955.78 GB	49,261.75...	8,694.03 GB	90.9%	20:00:00:0...		4G
	Switch 73	fc2		5,794.30 GB	4,924.47 GB	869.82 GB	9.1%	20:00:00:0...		4G
hdis300 (2)	scis300		0	32,213.36 GB	27,380.9...	4,832.41 ...	11.8%			4G
	scis300	fc2/14		16,193.21 GB	13,764.03...	2,429.17 GB	50.3%			4G
	scis300	fc2/15		16,020.15 GB	13,616.92...	2,403.24 GB	49.7%			4G
hdis200 (2)	scis200		0	24,496.01 GB	20,821.2...	3,674.78 ...	9.0%			4G
scis300 (2)	hdis300		0	17,410.65 GB	14,798.6...	2,612.01 ...	6.4%			4G
Switch 78 (1)	Switch 79	fc0	0	16,054.57 GB	13,646.1...	2,408.38 ...	5.9%	20:00:00:...		4G
hbrcd1 (1)	stbrcd1	fc15	0	10,183.44 GB	8,655.84 ...	1,527.60 ...	3.7%	20:00:00:...		2G
hbrcd2 (1)	stbrcd1	fc0	0	10,093.92 GB	8,579.75 ...	1,514.17 ...	3.7%	20:00:00:...		2G
F3brcd-b (1)	F3brcd-a	fc1	0	9,970.87 GB	8,475.14 ...	1,495.73 ...	3.7%	20:00:00:...		2G

4. Configure the storage array and fabric switches to enable all hosts to access the volume through multiple paths.
5. To verify that all hosts have access to the volume, select **Inventory > Hosts**.


6. Create a data store on the new volume and configure the preferred path from the hosts.
7. Select **Inventory > Virtual Machines** to display the **Virtual Machines** main view.
8. From the **Virtual Machines** main view, click on a virtual machine and click the **VMDKs** icon to show the virtual machine disks.
9. Verify that the data store was added to all hosts in the cluster.
10. Verify that there is no active path conflict (in the event of an active/passive storage array).

Determining if a VM is affecting host performance

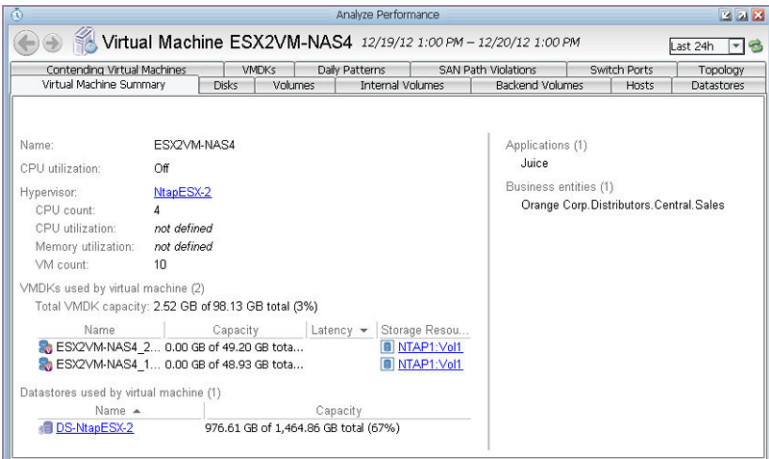
If a customer calls to complain that response time is unacceptable, you need to find the VM to determine whether it is responsible for the performance issue. For this example, the host name is NtapESX-2, the IP address is 192.1.168.39, and the name of the VM is vm-15.

Steps

1. Enter the name of the host, VM, or IP address into the Search facility.
In this example, the search results for the VM name show that only one VM matches the search criteria.



2. Double-click the selection in the search results that is most likely the VM you want to investigate.
The Virtual Machine Summary tab displays troubleshooting information.



3. Check to be certain that the correct business entity is represented in the summary and examine the data.

In this example, the summary shows that the Juice Sales application is running on this VM, and the CPU utilization is Off. That is an obvious problem, but what is the cause?

4. Based on the VM summary information, you might examine details on other tabs in the **Analyze** dialog box, or you might select the **Inventory > Virtual Machines** to view the VMs as a group and locate the one you have researched.

In this example using the Virtual Machine view, the state of the VM is "Shutting Down" and that VM has the smallest capacity of the VMs listed. This information reveals a problem with the way the Juice Sales application was set up since it is a High priority application.

Name	IP	Host Names	Host IPs	VM Capacity (GB)	VM Used Capacity (GB)	VM Memory (MB)	State	Application	Application Prio.
ESX1VM-NAS4	vm-5	NtapESX-1	192.1.168.38	343.77	10.42	6144	Unknown	Juice Sales	High
NtapESX-2 (10)									
ESX2VM-NAS2	vm-12	NtapESX-2	192.1.168.39	3,190.14	97.83	43,008	2048 Not Running	Juice	High
ESX2VM-NAS6	vm-13	NtapESX-2	192.1.168.39	245.57	7.93	4096	Resetting	Juice	High
ESX2VM-NAS5	vm-14	NtapESX-2	192.1.168.39	196.16	6.33	4096	Running	Juice	High
ESX2VM-NAS6	vm-8	NtapESX-2	192.1.168.39	393.14	10.62	6144	Running	Juice	High
ESX2VM-NAS4	vm-15	NtapESX-2	192.1.168.39	343.15	10.74	6144	Running	Juice	High
ESX2VM-NAS4	vm-9	NtapESX-2	192.1.168.39	98.13	2.52	2048	Shutting Down	Juice	High
ESX2VM-NAS5	vm-9	NtapESX-2	192.1.168.39	491.29	15.35	2048	Shutting Down	Juice	High
ESX2VM-NAS3	vm-16	NtapESX-2	192.1.168.39	441.50	12.98	4096	Standby	Juice	High
ESX2VM-NAS4	vm-10	NtapESX-2	192.1.168.39	343.38	10.09	4096	Standby	Juice	High
ESX2VM-NAS2	vm-17	NtapESX-2	192.1.168.39	441.87	14.73	6144	Unknown	Juice	High
ESX2VM-NAS3	vm-11	NtapESX-2	192.1.168.39	195.96	6.54	6144	Unknown	Juice	High
NtapESX-3 (18)									
ESX3VM-NAS7	vm-18	NtapESX-3	192.1.168.40	5,005.94	140.30	73,728	2048 Not Running	Juice	High
ESX3VM-NAS7	vm-24	NtapESX-3	192.1.168.40	343.15	11.07	2048	Not Running	Juice	High
ESX3VM-NAS7	vm-30	NtapESX-3	192.1.168.40	196.23	5.94	2048	Not Running	Juice	High
ESX3VM-NAS7	vm-30	NtapESX-3	192.1.168.40	392.24	10.32	2048	Not Running	Juice	High

5. At this point, you can respond to the inquiry.

Based on the research in this example, you conclude that the Juice Sales application needs to be moved to a larger capacity VM, and you inform the team and plan for that change.

Troubleshooting performance problems

When you receive calls reporting performance problems, OnCommand Insight provides the tools to use the reported facts along with your experience to track down the source of the problems.

While OnCommand Insight cannot describe the precise solution to a performance problem, it does give you tools enough to determine the best troubleshooting action based on information found in OnCommand Insight.

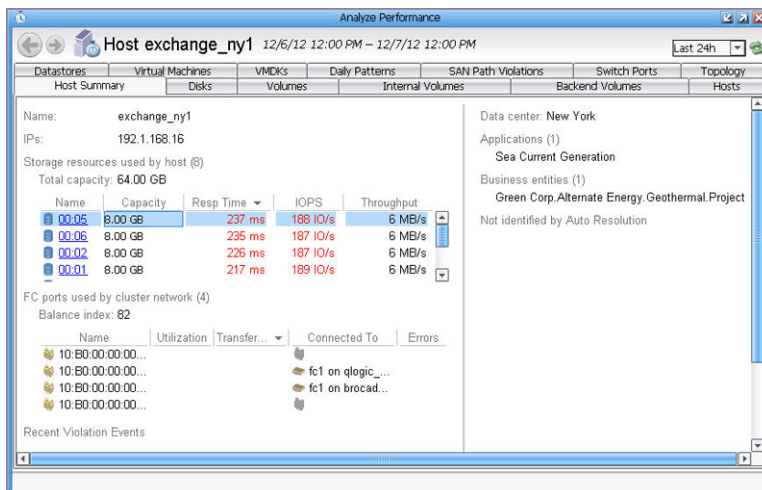
Note: To see a video demonstrating performance troubleshooting techniques, click this link to start the video on the NetApp Community: [Troubleshooting performance techniques video](#).

Troubleshooting slow client computers on your network

If you receive many calls from network users complaining that their computers are running slowly, you can troubleshoot host performance and identify the affected resources.

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Host Performance**.
2. Look for hosts with a high percentage of memory utilization. Select a host in this view.
3. Right-click and select **Analyze**.
4. In the **Host Summary** tab, check to be certain that the Data Center and Applications, shown on the right, are the ones with the reported slow client computers.



5. Click the link for the volume with a high response time or click the **Switch Ports** tab to examine a high transfer rate.
6. Determine the appropriate corrective action for your environment.

Related references

[Analyze dialog box](#) on page 70

[Host Summary tab](#) on page 75

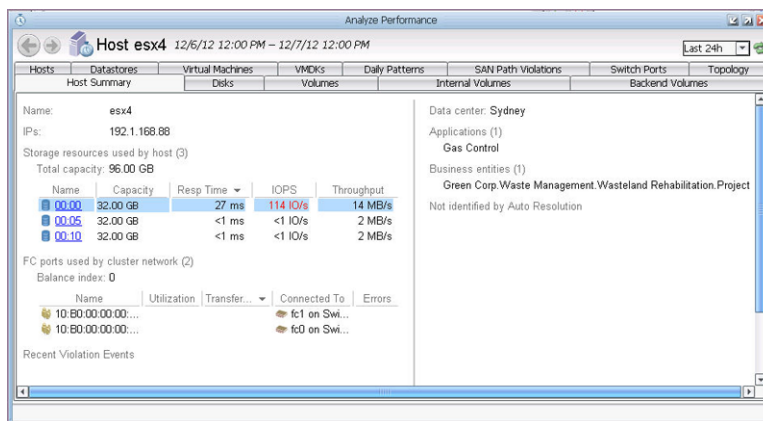
Troubleshooting poor application performance

Suppose the manager in a business unit called to complain that the application they are running is performing poorly. You can analyze the situation to identify the source of the problem and determine whether the poor performance is caused by the network or the application itself.

Steps

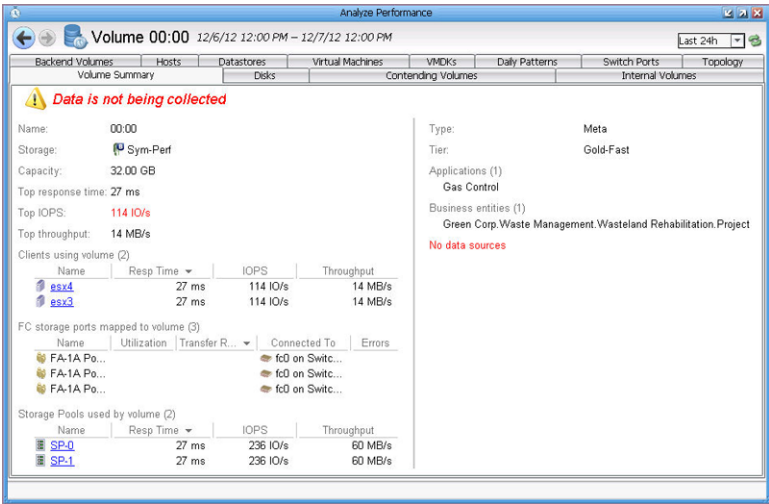
1. From the OnCommand Insight **Open** menu, select **Performance > Host Performance** to see the performance data for all of the hosts.
2. Select the server for the Data Center and Application that had a performance problem.
3. Right-click and select **Analyze**.
4. In the **Host Summary** tab, look for any high response times or a high transfer rates in the cluster FC network. Click the link to the volume with a high response time or transfer rate.

In this example, the storage resource has an IOPS of 114 IO/s that violates the performance threshold.



5. Click the link of the storage resource that is above the threshold.

6. In this example, the **Volume Summary** tab displays the resource details and indicates that there are no data sources. The error message and "No data sources" indicate that the data sources have stopped collecting data. This failure might not be the source of the performance problem, but rather the result of that problem.



7. At this point, the technical team must explore the data source failure.

Related references

[Host Performance view](#) on page 93
[Volume Summary tab](#) on page 79
[Host Summary tab](#) on page 75

Analyzing traffic flow and identifying congestion

Smooth traffic flow and preventing congestion on a storage network device or across devices is a primary concern of administrators. Storage network congestion often impacts service levels and might cause application failures.

Storage network traffic is unpredictable. Depending on an application's demands, the amount of data flowing from hosts through switches, ISLs, and storage arrays can fluctuate greatly. While some applications send and receive less data over time, others send and receive vastly increasing amounts.

Hosts are a crucial part of the storage network because they house the applications that perform vital business operations. When a host becomes congested, business applications experience degraded performance or fail. You can use these OnCommand Insight tools to analyze traffic flow and pinpoint areas of congestion to prevent many problems:

- Switch Port Performance
- Analyze
- Analyze Contention
- Switch Port Performance Alerts
- Daily Traffic Pattern chart
- Host Performance

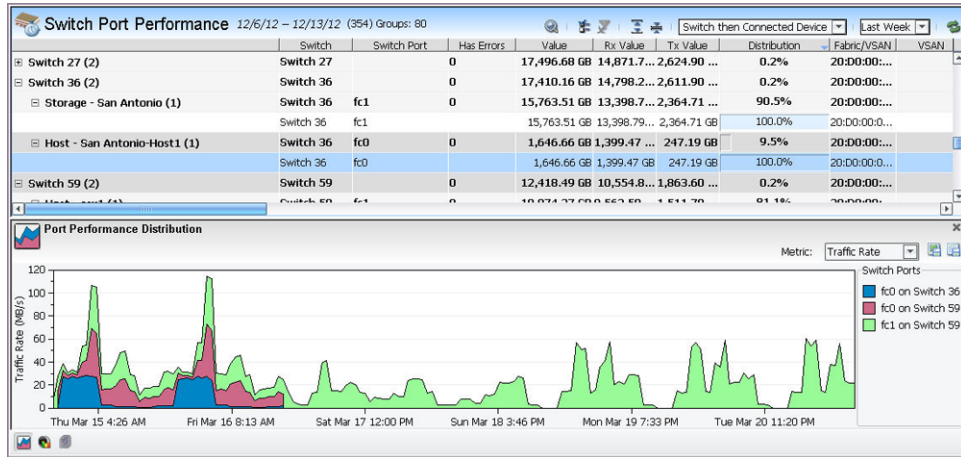
Analyzing traffic flow

You can use examine switch port performance and analyze contention to see the traffic flow. The detailed charts illustrate when there are traffic pattern spikes or troughs during a single day. You can select multiple ports and see the collective impact of traffic on the switch or check individual ports for periods of high or low traffic.

Steps

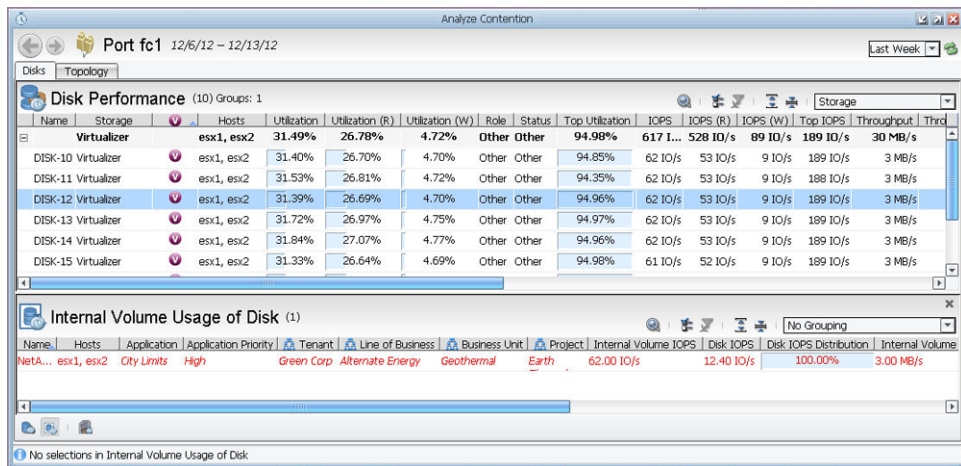
1. Select **Performance > Switch Port Performance**.
2. Select the switch ports that you want to analyze and expand the list to see the Distribution percentages and if there are any errors.
3. To examine a detailed chart of the traffic flow, click the **Port Performance Distribution** icon and adjust the **Metric** to show Traffic rate.

In this example, the graphic comparison of the switch ports shows that two ports have spikes in their traffic flow.



4. Select the switch ports with high traffic from the list in the **Switch Port Performance** view.
5. Right-click and select **Analyze Contention**.
6. Scroll across the **Disk Performance** list to examine the utilization percentages.
7. Select an item of concern and click one of the icons to display the details.

In this example, the internal volume with a high top utilization percentage is running the high priority City Limits application.



8. At this point, you might select **Performance > Application Performance** or return to the **Switch Port Performance** view to analyze a different switch.

Related references

[Analyze Contention dialog box](#) on page 68

Identifying a congested host

Many times you might suspect that a congested host is the result of a slow switch or over-utilized volume, when in reality, the host is congested because other hosts are over-utilizing the same storage port. You can identify a congested host and the cause of the congestion by noting errors occurring on all hosts. Hosts receiving errors are likely experiencing congestion.

About this task

If you already know the name of the problem host, skip this procedure.

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Switch Port Performance**.
2. From the menus in the upper right corner, select **Connected Device Type then Name** and **Last Week**.

This displays all connected devices including storage devices and the amount of data traffic in Megabytes per storage device.

3. Expand the list of Hosts.
4. To examine the errors on each host in the **Has Errors** column, select a host with a high number of errors and a high amount of data traffic as a potential candidate.
5. To copy the host name, click the host name in the first column and press **Ctrl+C** to copy the host name into your computer Clipboard.

After you finish

You need the Connected Device Name of the potential congested host in your Clipboard to start the analysis of the host congestion.

Identifying the cause of the congested host

To identify the cause of a congested host, you need to identify the storage ports that the suspected hosts access as well as all other devices that access the same storage ports as the congested host.

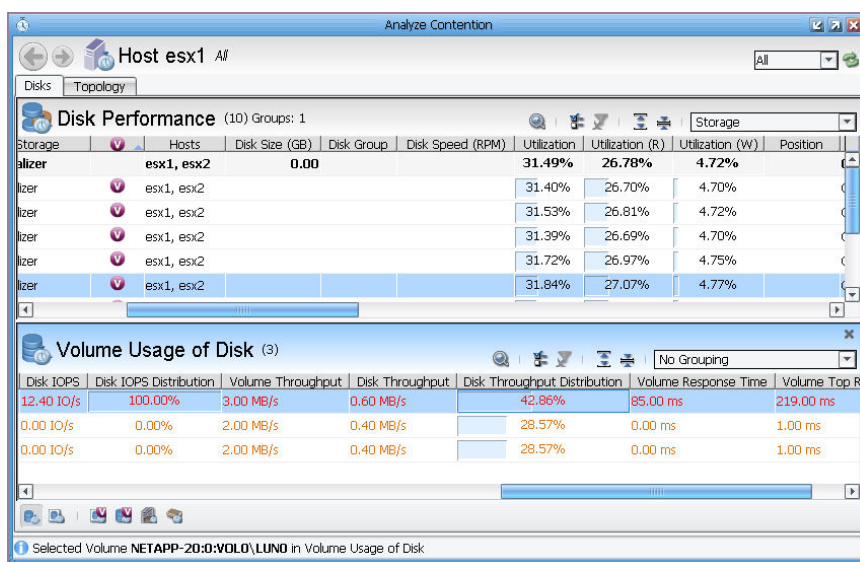
Before you begin

Before starting this procedure, copy the Device Name of the potential congested host into the Windows Clipboard.

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Host Performance**.
2. Sort the displayed hosts by name or click the Filter icon in the Name column and paste the name of the host by pressing **Ctrl+V** and pressing **Enter**.
3. Right-click the host and select **Analyze Contention**.
4. For the list of Disks, examine the disk performance and utilization information.
5. If you need the capacity information for a disk, select it, right-click, and select **Show Storage Pools**.
6. To display the amount of data flowing through the selected disk, in the **Analyze Contention** dialog box click the **Volume Usage of Disk** icon.

In this example, the high Disk IOPS Distribution, Disk Throughput Distribution, and response times reveal the contention problem.



Related references

[Analyze Contention dialog box](#) on page 68

Analyzing congested storage ports

When a storage port becomes congested, hosts experience degraded performance, and application performance also suffers. You can locate congested storage ports, analyze the cause of the congestion, identify SAN devices most likely responsible for the congestion, and solve the problem.

The two most common causes of congestion are a high amount of host data traffic (from a busy application) accessing a volume through the storage port or a malfunctioning switch port.

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Switch Port Performance**.
2. From the three drop-down lists in the upper right corner, select **Connected Device Type then Name, Last Week, and Traffic**.
3. Expand the **Storage** group to locate a potential candidate with a high number of errors and a high amount of data traffic.
4. Expand the storage array and write down the name of the storage array ports or select an individual port.

To copy the storage name, click the storage name in the **Connected Device Name** column and press **Ctrl+C**.

5. From the OnCommand Insight **Open** menu, select **Inventory > Storage Arrays**.
6. To sort the storage devices by name, click the **Name** column heading.
7. Paste the copied name of the problem storage array into the filter by pressing **Ctrl+V**. Press Enter.
8. Click the **FC Ports** icon.
9. Right-click on the storage array port and select the **Show Daily Traffic Pattern** option to display the Histogram View for the port.

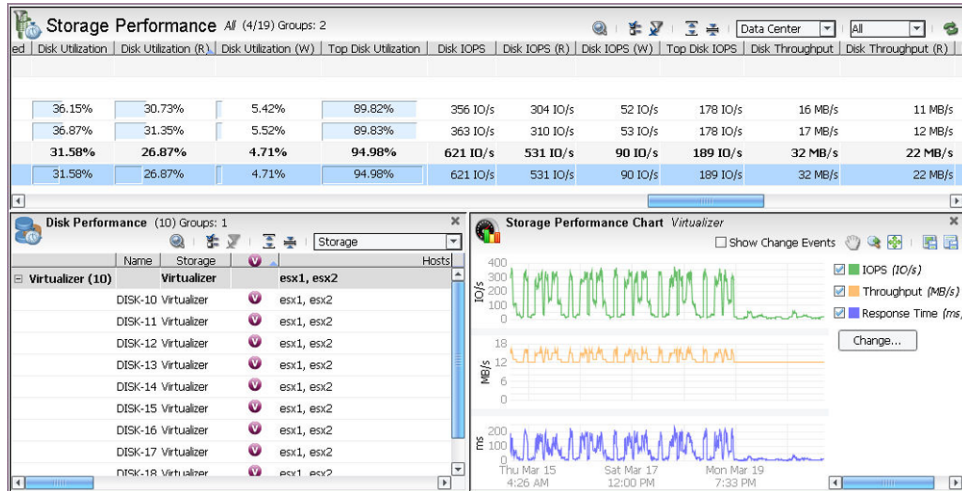
Review the list of all the devices (by type) that access the congested storage port. The congested storage port also appears. The amount of data flowing through these devices to all storage arrays is displayed. The host displaying the most data traffic is the probable cause of the storage port's congestion.

Analyzing virtual storage performance

You can identify the source of performance problems on the virtualizer, on the virtual disks used by the virtualizer, and on the virtual volumes and internal volumes created by the virtualizer.

Steps

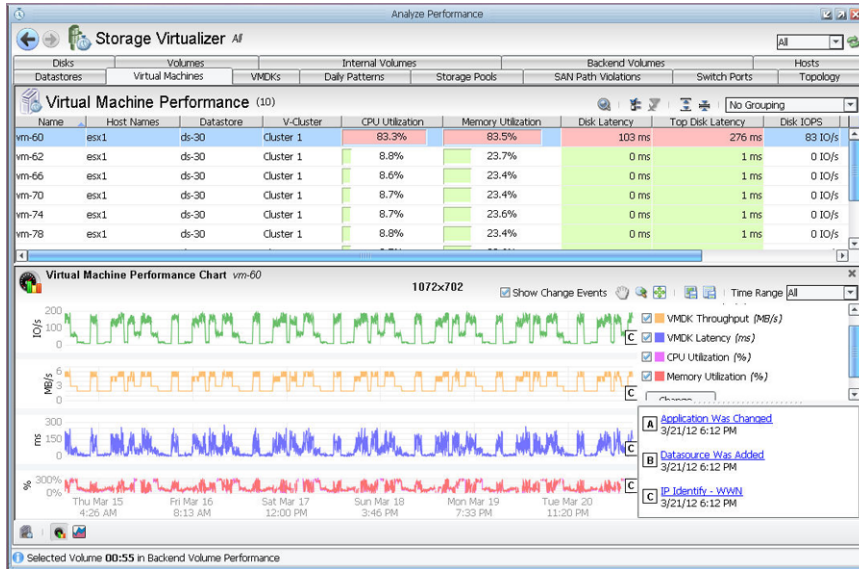
1. From the OnCommand Insight **Open** menu, select **Performance > Storage Performance**.
2. Filter the columns to show only the virtualized storage in the view.
3. Examine the data for any anomalies particularly in the utilization columns.



4. To analyze the performance of a selected virtualized storage, right-click, and select **Analyze**.
5. Click the **Analyze** tabs to display different data:
 - Disk Performance view: Shows the performance of the virtual disks used by the virtualizer.
 - Volume Performance view: Shows the performance of the virtual volumes created by the virtualizer.
 - Internal Volume Performance detail view: Shows the performance of the virtual internal volumes created by the virtualizer.
 - Backend Volumes Performance detail view: Shows the performance of all the backend volumes used by the virtualizer.
6. To see a chart summarizing the distribution of performance metrics throughout a day for all the disks, volumes, internal volumes, and backend volumes, click the **Daily Patterns** tab.
7. For the Disks, Virtual Machines, Volumes, Internal Volumes, and Backend Volumes tabs, click the icons to display the following for the selected rows:
 - Performance chart showing performance metrics over time, such as IOPS and throughput and also has markers for change events.
 - Distribution chart showing the distribution of performance metrics of all selections over time.

In this example, the high utilization and latency problems shown on the Virtual Machines tab are selected, and the Virtual Machine Performance Chart shows that three change events occurred at the same time to affect the performance:

- Application was changed.
- Datasource was added.
- Device Created -WWN.



- Click the links in the performance chart to display the Changes for each event and determine which events caused the problems.

Balancing resources

A key goal for the storage administrator is to use capacity in a proactive way to avoid problems like contention for resources. Administrators need to predict and prevent these problems. You can use OnCommand Insight to analyze the network traffic on your storage network as well as the distribution of that traffic on all devices and ports and determine how to balance the data flow among them.

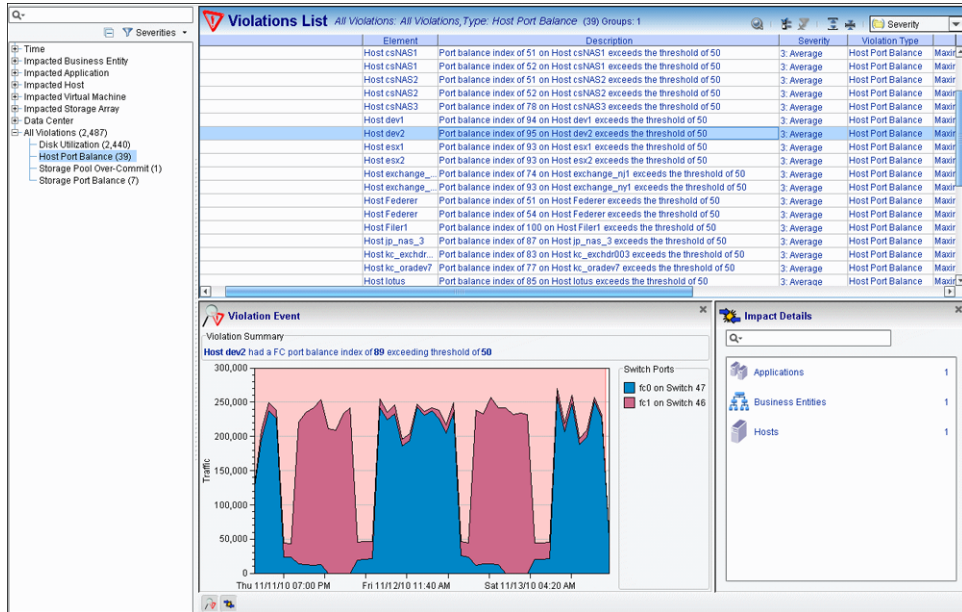
Balancing the data flow means extracting more utilization out of the storage network without impacting application performance.

Analyzing port balance violations

When examining port balance violations, you need to determine whether the violations indicate ongoing problems or temporary conditions.

Steps

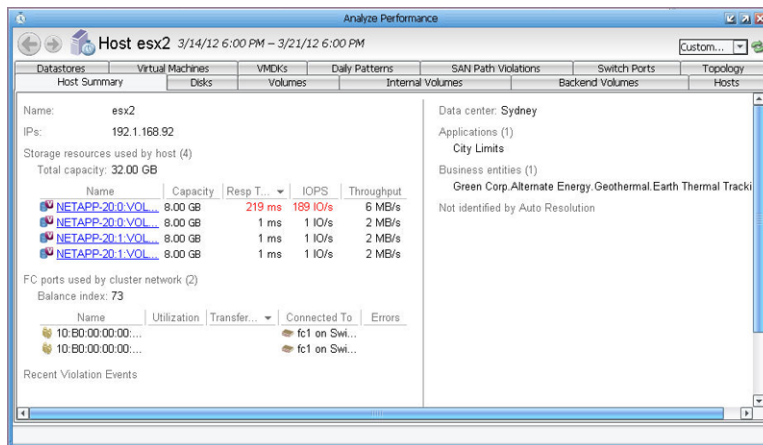
1. To examine the list of host port balance violations, select **Assurance > Violations Browser**.
2. In the browser tree, expand **All Violations** and select **Host Port Balance**.
3. Click the **Violation Event** icon to see an overview of the port traffic balance over time.
4. Click the **Impact Details** view to see the resources impacted by the port balance violation.



5. Select one violation of interest, right-click, and select the **Analyze** option.

Example

In this example, the Host Summary tab shows that one volume has a response time beyond the threshold.



6. Click the link to the display more details for the volume with the high response time.

7. Select the **Contending Volumes** and **Daily Patterns** tabs to determine if the out-of-balance condition is a trend or a special case.

After you finish

If the port traffic over the device is permitted to be out-of-balance, you can modify the policy on the device to increase the balance index threshold or disable the policy on the device completely to eliminate the violation.

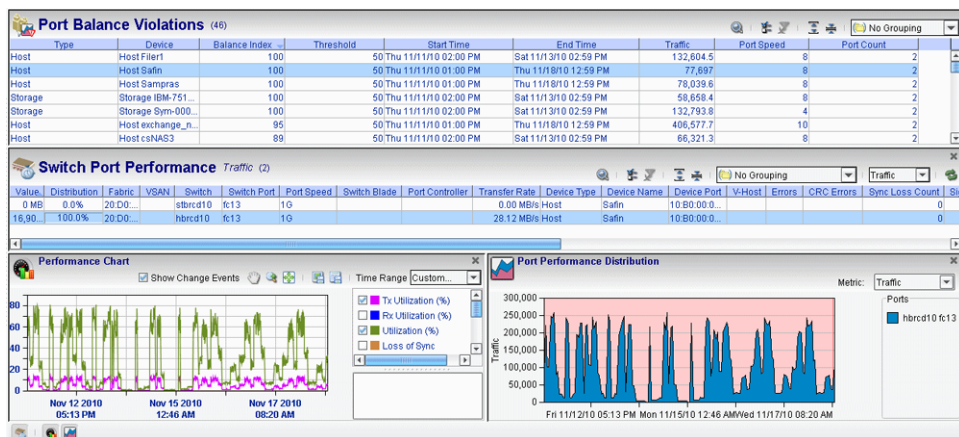
Identifying out-of-balance network traffic across host ports

If the ports of a host are connected and active but some are transmitting or receiving most of the data while others are transmitting or receiving very little, this is a symptom of out-of-balance ports.

Steps

1. From the OnCommand Insight **Open** menu, select **Assurance > Port Balance Violations**.
2. Click the **Balance Index** column to change the sort order of the Hosts by the Balance Index value.
3. Select a host with a high Balance Index compared to its threshold.
4. Click the **Switch Port Performance** icon and select a line describing a switch that interests you.
5. Click the **Performance Chart** and **Port Performance Distribution** icons.

Review the detailed information see if a host is sending much more data through one switch than another. If this is the case, one switch may be over-utilized, which could result in slower applications, or if the switch carrying the data fails or is taken offline, the host will experience severe performance degradation.



Identifying resources for storage allocation

After you know how much data you have and where and how much bandwidth current individual network paths (switches, ports) provide, you can decide how much data can be allocated to existing devices and the path through which the data can flow. Then you know if new hardware is necessary, and where on your network new hardware is most needed.

About this task

To identify resources for new storage allocation, follow these general steps:

Steps

1. Identify a storage array that meets the requirements including available space, bandwidth, and tier.
2. Select storage ports based on performance and connection.
3. Identify volumes with potential capacity based on their array, speed, and availability.
4. Select switch ports based on the amount of available bandwidth.
5. Identify Tier 1 to Tier 2 storage migration candidates.

Analyzing storage port balance violations

When examining storage port balance violations, you need to determine if the violations indicate legitimate problems or not.

Steps

1. To examine the list of storage port balance violations, select **Assurance > Violations Browser**.
2. In the browser tree, expand **All Violations** and select **Storage Port Balance**.
3. Select one violation of interest, right-click, and select the **Analyze Violation** option.
4. In the **Port Balance Violations** main view, select the violation to display switch port performance data and the performance charts.

After you finish

If the device ports should be balanced but are not, change the device configuration to resolve the issue. However, if the device port traffic is permitted to be out-of-balance, you can modify the policy on the device to increase the balance index threshold or disable the policy on the device completely to eliminate the violation.

Identifying an available storage array

As a first step in identifying resources for reallocating storage, you need to identify a storage array that is a good capacity candidate.

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Switch Port Performance**.
2. To display all devices including storage devices, the amount of data traffic, and amount of data traffic per storage device, select **Connected Device Type then Name, Last Week** and **Traffic** from the menus.
3. If your storage array must be in a certain location or of a certain type, click the title bar of the storage array to open the filter and type letters or words to display the appropriate arrays.
4. Select a storage array with a low distribution percentage.

A low distribution percentage indicates that the storage device has been used very little over the past week and has neither received nor sent much data, making it a good resource candidate for new storage allocation. The array could also be used to house data for an existing over-utilized storage array.

Selecting storage ports

The second step in the resource allocation is to identify the storage port(s) that are good candidates for your capacity needs.

Steps

1. After identifying the best storage array in the Switch Port Performance view, expand the storage device to examine the list of ports.
2. Select storage ports with a low distribution percentage and low value.

A low value in both categories is optimal. This indicates that the storage port has been used very little, making it a good candidate to mask volumes to the port.

Identifying volumes with potential capacity

After the best storage array and storage port are identified, select the best volume for your capacity needs.

Steps

1. From the OnCommand Insight **Open** menu, select **Inventory > Storage Arrays**.
2. Select the storage array chosen earlier. You can sort or filter by name.
3. Click the **Masking** and **Volumes** icons.
4. Review all the available volumes on the selected storage array.

5. In the Masking view, expand the No Mapping and/or No Masking categories and select the volume that meets your requirements for available space and capacity.

The No Mapping and No Masking views display only those volumes for the selected storage array that have not yet been mapped or masked, respectively, in the network.

Identifying Tier 1 to Tier 2 data migration candidates

To use your storage network as efficiently as possible, applications should be matched to devices that deliver appropriate services. High-traffic applications using critical data need Tier 1 storage. Lower-traffic applications can safely use lower-tier, cheaper storage devices.

About this task

Applications that have requested and are assigned Tier 1 status sometimes use only a fraction of Tier 1 performance and can be safely migrated to cheaper storage arrays, freeing expensive storage capacity to make better use of existing storage devices and avoid the unnecessary purchase of additional (and expensive) Tier 1 storage.

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Switch Port Performance**.

2. Select **BU** then **Application** then **Connected Device** and **Last Week** from the menus.

This displays all business entities, the applications they use, and the devices those applications reside on. This view shows the amount of data traffic in Megabytes that has traveled through each device over the past week.

3. To sort the storage arrays experiencing the most data traffic in descending order, click the **Distribution** column heading.
4. Expand a business entity and then the applications that are experiencing low levels of data traffic.

If the applications showing little data through their respective hosts reside on a Tier 1 storage array, they can safely be migrated to a lower-tier array. In the current view, you can also expand the **Name** column to show the individual hosts and to see the amount of data flowing to all its ports.

5. To copy the name of a selected host, highlight the line for that host and double-click on the name in the **Connected Device Name** column. Press **Ctrl+C**.

You can use this text in the filter to pinpoint the host in the next view.

6. From the OnCommand Insight **Open** menu, select **Inventory > Storage Arrays**.

If the storage arrays for that host are Tier 1 arrays, it is advantageous to move one or more applications on the host to cheaper, lower-tier arrays.

Balancing data flow across storage array ports

If large amounts of data are flowing through a small number of storage array ports, the array cannot operate efficiently. To maximize the performance of your storage arrays, ensure that the data is distributed equally across all storage array ports.

About this task

Redirecting data traffic going to the ports on the array through one of the non-utilized (or in most cases, under-utilized) storage ports will alleviate bottlenecks. Additionally, a large amount of non-utilized ports makes this array a good candidate to house data for new applications or to house data from other storage arrays that are experiencing very high amounts of data traffic.

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Switch Port Performance**.

Be certain that **Connected Device Type then Name** and the time period are selected in the drop-down lists.

2. Expand the storage list in the first column. Select a storage array and expand it.
3. To sort storage array data distribution from high to low, click the **Distribution** column.
This displays the data traffic flow across the storage ports on the array in descending order.
4. Highlight one or more storage ports of interest and click the **Port Performance Distribution** icon.
5. Identify any ports that are receiving the majority of the traffic while others are receiving very little.

Balancing data flow across ISLs

Inter-Switch links (ISLs) connect high-speed switches. If large amounts of data are flowing through only one ISL port while other ISL ports experience small amounts of data flow, the ISL is imbalanced and operating inefficiently. To maximize the performance of your ISLs, ensure the data is distributed equally across all ISL ports.

Steps

1. From the OnCommand Insight **Open** menu, select **Performance > Switch Port Performance**.

Be certain that **Connected Device Type then Name** and **Last Week** are selected in the drop-down lists.

2. To sort the data distribution from high to low, click the **Distribution** column.

This displays the data traffic flow across the switch ports in descending order.

3. Select an individual switch and expand it.

You can see the traffic moving between ISLs on the expanded switch. The switch under **Connected Device Name** is the switch that is expanded in the far-left column. The second switch is identified under the heading **Switch**.

4. Check to see if any links are carrying more data than other links.

It is very likely that this imbalance is the result of a misconfiguration or another problem and should be corrected as soon as possible.

Allocating capacity to a new host

In order to implement the allocation of capacity to a new host, the administrator needs to look at the array and select an array that best fits the request. Using filters, view all of the arrays that exist in the data center (using business entities or annotations) and then select an array with the lowest disk utilization.

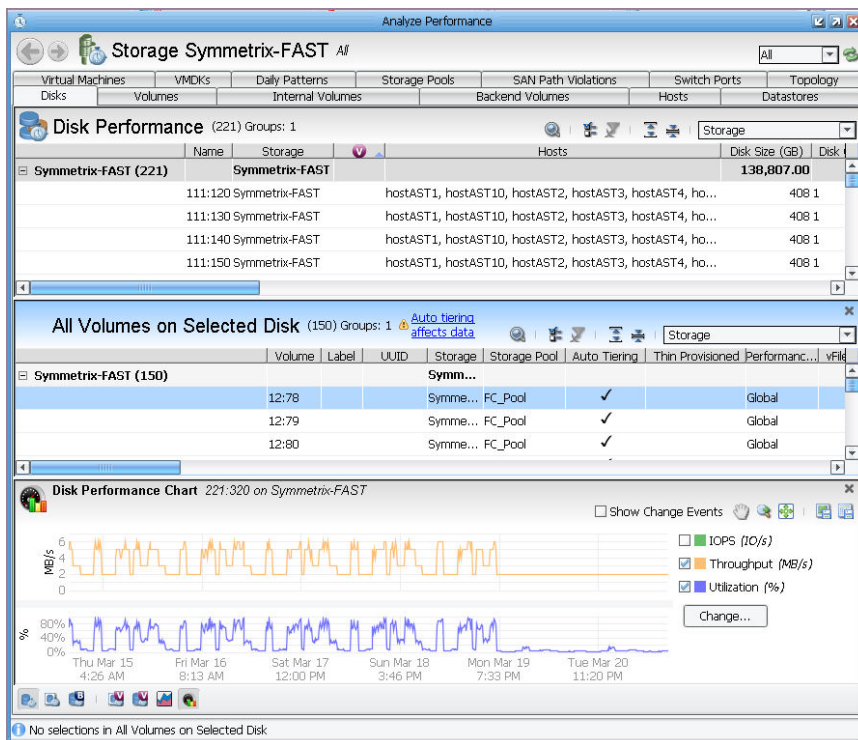
Steps

1. Open **Performance > Storage Performance**.
2. Select one or more arrays that you believe might be under-utilized based on the Disk Utilization percentages in the Storage Performance view.
3. Click the **Disk Performance** detail view icon to show more information about the disk utilization for the selected arrays.

In this example, the information is grouped by the Data Center annotation.

Storage Performance All (19) Groups: 4										
Name	Technology	Capacity (GB)	Raw Capacity (GB)	Vendor	Family	Model	Volume Response Time	Volume Response Time (R)	Volume Response Time (W)	Volume Response Time (M)
Sym-000050074300328		3,100.00	6,000.00 EMC	Sym...	Sym...		92 ms		92 ms	
CK600_29_CK00351029329	FC	7,448.00	7,448.00 EMC	Clarion	CX600		0 ms		0 ms	
Zurich (6)		126,712.54	185,956.00 EMC, NetApp							
APM000910025			EMC	Celer...	Celer...					
Symmetrix-FAST		113,476.54	162,868.00 EMC	Sym...	Sym...		84 ms		82 ms	
V-Series-Zurich2	FC/SCSI, Cl...	64.00	NetApp	V-Series V-Series			68 ms		67 ms	
V-Series-Zurich1	FC/SCSI, Cl...	64.00	NetApp	V-Series V-Series			70 ms		69 ms	
Sym-000050074300324		3,000.00	5,000.00 EMC	Sym...	Sym...		78 ms		76 ms	
CK600_26_CK00351029326	FC	10,108.00	18,088.00 EMC	Clarion	CX600		0 ms		0 ms	
Sydney (4)		17,076.00	30,864.00 EMC, NetApp				83 ms		81 ms	
New York (4)		21,118.00	30,108.00 EMC, HDS, HP				86 ms		85 ms	
Disk Performance (221) Groups: 1										
Name	Storage	Hosts	Disk Size (GB)	Disk Group	Disk Speed (RPM)	Utilization	Utilization (R)	Utilization (W)	Position	Role
Symmetrix-FAST (221)			138,807.00			31.58%	26.85%	4.73%		Data
111:120 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.64%	26.90%	4.74%		Data	
111:130 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.57%	26.84%	4.73%		Data	
111:140 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.64%	26.90%	4.74%		Data	
111:150 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.58%	26.85%	4.73%		Data	
111:160 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.93%	27.15%	4.78%		Data	
111:170 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.76%	27.00%	4.76%		Data	
111:180 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.26%	26.58%	4.68%		Data	
111:190 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.65%	26.91%	4.74%		Data	
111:200 Symmetrix-FAST	hostAST1, hostAST10, hostAST2,...	408 1	15000		31.80%	27.04%	4.76%		Data	

4. If necessary, select different arrays in the Storage Performance view to analyze.
5. After you have selected the arrays that are the most under-utilized, right-click those array descriptions in the Storage Performance view.
6. Select the **Analyze** option. The Disks tab is open in the **Analyze** dialog box.
7. Select a disk in the Disk Performance list that you want to investigate.
8. For the selected disk, click the **Volume Performance** icon to display the **All Volumes on Selected Disk** detail view. This view reveals the volumes that reside on the disk.
9. You might also want to display the **Performance Chart** for that volume, as shown in this example.



10. In the **All Volumes on Selected Disk** detail view, change the data grouping to **Storage** and select the volume to use as the new target volume for the capacity requirement.

Allocating capacity to an existing host

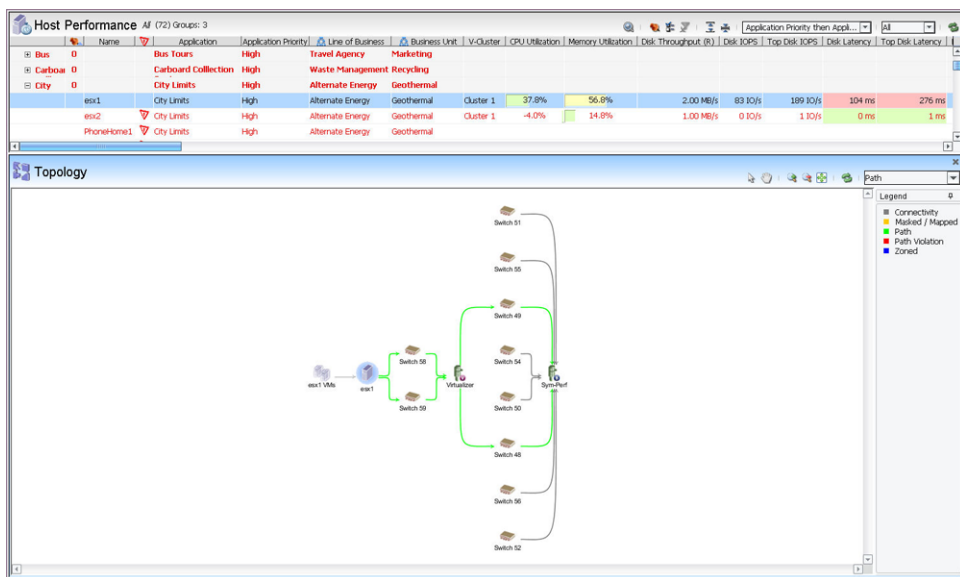
Usually when a host is connected to an array, the storage administrator tries to allocate the next volume for that host from the same array unless the array is hot (in other words, unless devices are already in heavy use on a regular basis or at a critical threshold).

Before you begin

To perform the allocation, the administrator must first locate the arrays that are connected to the required host.

Steps

1. Select **Performance > Host Performance**.
2. Select a grouping for the data. This example is grouped by Application Priority then Application.
3. Locate the host that needs to be allocated more capacity.
4. Click the **Topology** icon to determine which arrays are connected to the selected host.
5. To see all of the devices, click on an item in the topology and select the **Add All Connected Devices** option.



6. Select an array in the topology view and click the **Disk Performance** icon.
7. Identify the disk with the lowest utilization in the Disk Performance detail view.

8. Right-click on that disk and select the **Analyze** option. The Disk Performance view shows only data for the chosen disk.
9. Sort this data using the grouping data for **Storage**.
10. Select a volume from the list of volumes. That volume might be a good choice for the volume for the capacity requirement.

When selecting a new volume to allocate on an existing host, the storage administrator can also use the Topology view to see the storage that the host is already using.

Related references

[*Disk Performance view*](#) on page 89

Performance reports

You can generate two reports from your OnCommand Insight performance data for utilization and identifying orphaned volumes. The list of the top 10 most and least utilized hosts, under utilized volumes, ISL Balance index, and unbalanced devices helps you identify the most utilized hosts, ISLs and arrays.

One challenge to improve utilization in a storage network requires the identification of orphaned volumes. A storage administrator wants to reduce capital expenditures by identifying under utilized volumes and then reclaiming or consolidating them. The storage administrator needs to determine which volumes have not received significant volume throughput and thus are candidates for reclamation.

Note: You might also want to use the performance reports available from the Reporting Portal. See the *OnCommand Reporting Guide* for more information.

Gathering performance information

You can use the Perform Report to explain the deferment of capital expenditures, apply non-utilized resources to new projects or applications, identify resources for consolidation, or justify a tiering strategy.

Before you begin

You must have a Perform license.

Steps

1. From the OnCommand Insight portal, select **Reports > Insight Perform Report**.
2. Set these parameters to generate the report.
 - **Report on:** select Last hour, Last day, or Last week.
 - **The number of utilized/underutilized Hosts, Arrays and ISLs:** enter a number.
 - **Unutilized Volumes Traffic Threshold (MB):** enter a number.
 - **Unbalanced ISLs Index Threshold:** enter a number.
3. Click **Generate Report**.
4. Use the report tabs in the spreadsheet to examine this utilization information:
 - **Summary** shows the start and end timestamp for the measurements taken, the performance threshold you have set for the host and the ISL (Inter-Switch Link) Balance Index.
 - **Top 10 most and least utilized hosts** allows you to assess the effectiveness of load balancing in the storage network. Highly utilized storage devices experience near-threshold data traffic, indicating the possible need to relocate applications to less utilized storage arrays. The Most

Utilized list also helps determine whether redirecting data traffic going to ports of specific arrays through under-utilized storage ports will alleviate the potential for bottlenecks. A low distribution percentage in the Least Utilized list indicates that the storage device has been used very little over the designated time period and has neither received nor sent much data, making it a good candidate for new storage allocation.

- **Unutilized Volumes** list summarizes the storage resources in the SAN that are not being utilized at all and which are available.
- **ISL Balance Index** is the Inter-Switch Link (ISL) Balance Index, a calculated Standard Deviation, that measures how balanced the load is across switch ports.
- **Unbalanced Devices** is the balance index, a calculated Standard Deviation. It is a measurement of how balanced the load is across switch ports. The further away from 1 that the ISL Balance Index is, the more likely the host load is going to be imbalanced across switches. A high value (above 50) signifies a ratio problem between the host ports; that is, some ports are experiencing a large amount of traffic and others a very low amount, indicating a potential problem. This demonstrates the need to balance data flow across ISLs.

Reporting array performance for orphaned volumes and tiers

Use the Array Performance Report to identify orphaned volumes. The tiered volume section helps the administrator manage tiering and meet capacity demands in the storage network.

Before you begin

You must have a Perform license.

About this task

Correcting inefficient utilization in a storage network requires the identification of orphaned volumes. A storage administrator wants to reduce capital expenditures by identifying under-utilized volumes and then reclaiming or consolidating them. The storage administrator needs to determine which volumes have not received significant throughput and are candidates for reclamation.

Steps

1. From the OnCommand Insight portal, select **Reports > Insight Perform Array Performance Report**.
2. Set the parameters to generate the report, which is based on the throughput (I/O per sec). Throughput is based on the interval of Last Hour, Last Day, or Last Week.
3. Click **Generate Report**.

The tabbed spreadsheet includes these sections:

- **Summary** of orphaned volumes displays the storage information and Raw Capacity in GB for each.

- **Orphaned Volumes** provides the specific breakdown of orphaned volume data. You can identify orphaned volumes measured by throughput.
- **Tiered Volumes** provides the specific breakdown of tiering data, with a throughput threshold of >100 I/Os. The report shows volume capacity (GB), raw capacity (GB) and total I/Os.

Performance reference

This reference material provides detailed descriptions of the main views and detail views for the **Performance** features.

This reference lists the items in alphabetical order. The view descriptions are displayed when the views are open and you press the F1 key or move the pointing hand icon to the view and click.

All of the available columns in each view are defined. However, the number of columns displayed initially in any view is limited to a default set. You can change the columns displayed and the order of these columns to suit your needs. See the "Showing or hiding columns in tables" and "Changing column order" instructions in the *OnCommand Insight Inventory User Guide*.

Automatic storage tiering in Inventory and Performance views

Data in views is affected by the use of automatic storage tiering.

If a host uses a volume that is configured for automated storage tiering (for example, FAST VP), OnCommand Insight does not show all of the storage pools for the volume. You see only the primary, native pool for that volume. If, as part of the automatic storage tiering, a volume moves to another storage pool, OnCommand Insight does not display the new relationship.

The following views are impacted by the use of automatic storage tiering (for example, FAST VP):

- **Analyze Contention:** The Disks tab does not show all the disks, and information about the relationship between disks and volumes is incomplete.
- **Analyze:** The Disks tab does not show the full list of disks; it shows only the primary storage pool disks. The Backend Volumes tab does not show the full list of backend volumes; it shows only the backend volumes of the primary storage pool. The Volumes and Backend Volumes center views of the Disks tab do not show all the relevant volumes. The Disks center view of the Volumes tab does not show all the relevant disks. The Hosts tab does not show all the hosts using a disk. The Daily Patterns tab does not show complete information when **Disk** or **Backend Volume** is selected for Performance data.
- **Backend Volumes and Backend Volume Performance detail views:** Do not show all the relevant backend volumes.
- **Disks and Disk Performance detail views:** Do not show all the relevant disks.
- **Disk metric columns of the Host Performance view:** Because roll-up values do not include all the relevant disks, these columns contain N/A values.
- **Host Performance Chart view and Host Performance Distribution view:** No data is shown if automatic storage tiering would display incorrect data.
- **Storage Pools and Storage Pool Performance detail views:** Do not show all the relevant storage pools.

- Virtual Storage detail view: Does not show all the relevant virtual disks and backend volumes.
- Volume Members By Disk dialog box: Does not show all the relevant volume members.

Related references

[Analyze Contention dialog box](#) on page 68

[Analyze dialog box](#) on page 70

[Disk Performance view](#) on page 89

[Performance Distribution view](#) on page 103

Analyze Contention dialog box

The Analyze Contention dialog box allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are contending with each other for performance on the disk, a factor that can cause poor performance.

Navigation

From the OnCommand Insight Open menu, select any Performance view or the Hosts, Virtual Machines, Datastores, Paths, or Storage Arrays views in Inventory. Right-click a device in the view and select the Analyze Contention option.

Tabs

Select either of these tabs to display additional information about the selected items in the view.

- Disks
- Topology

Column descriptions

The Analyze Contention dialog box shows all disks connected to the selected device with one row per disk.

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	Name of device.
Storage	Name of the storage array.
<i>array virtualization type icon</i>	Indicates virtualization and whether the storage acts as a virtualizer device (V) or as backend storage (B).
Hosts	Hosts making use of the disk.

Disk Size (GB)	The physical disk capacity, in gigabytes. This differs from usable capacity when technologies such as RAID are used, where some of the raw capacity is used for protection purposes.
Disk Group	Name of the set of disks that share a common configuration. A disk group contains volumes that must use disks within that disk group.
Disk Speed (RPM)	How fast the disk spins, in RPM. (This value is not applicable to SSD disks.)
Position	Where the disk is located (for example, as in Symmetrix: DirectorID-075, Slot 38).
Role	Role of the disk, as obtained from the vendor. Values include data, faulty, parity, double parity, mixed, spare, unallocated, or other. For spare disks, "spare" is displayed.
Status	Status of the disk. The following values could appear: Normal, Copying, Failed, Partner, Pending, Questionable, Reconstructing, Zeroing, Other.
Vendor	Name of disk vendor.
Type	Type of disk. One of the following values appears: ATA, FATA, FC, SAS, SATA, SATA2, SATA3, SSD, Other.
Model	Name of disk model.
Serial Number	Serial number of the disk, if the data source supports reporting it.
Top Utilization	The maximum percentage of the total pre-cache Service Time that can be used for Read and Write requests out of the selected time range.
IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).
IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum number of I/O service requests that can pass through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Throughput	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Throughput (R&W)	Rate at which Read or Write data is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Throughput	Maximum rate at which data can be received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Accessed	"Yes" indicates that the resource was accessed in the selected time range.
Application	Name of application.

Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.

Related concepts

[Automatic storage tiering in Inventory and Performance views](#) on page 67

Related tasks

[Identifying the cause of the congested host](#) on page 48

[Analyzing traffic flow](#) on page 46

Analyze dialog box

The Analyze dialog box can be accessed from many different views to identify potential problems by isolating the location of contention and the source of delays and by comparing performance data for different time periods to note changes. The Perform license is required for this feature.

Navigation

- Highlight one or more items in one of the performance views. Select the **Action > Analyze** options.
- Select one of the performance related items in the search results.
- From the OnCommand Insight Open menu, select **Performance** and any of the options. Select one or more items in the selected view and right-click. Select the Analyze option.
- To return to a previous view of the dialog box, click the left arrow in the upper left corner. To move forward in the tab views, click the smaller right arrow.
- To display a breadcrumb list of the previously displayed views, click and hold the arrow with the blue background (active).

Tabs

Select any of these tabs to display additional information about the selected items in the view. You might also want to select items on these tabs and display detail views with even more specific information. These tabs display the performance data for the selected items.

- <resource type> **Summary** is the dashboard for research into a selected performance problem with links to more information.
- Disks
- Volumes
- Internal Volumes
- Backend Volumes

- Switch Ports
- Storage Pools
- Hosts
- Datastores
- Virtual Machines
- VMDKs
- **Daily Patterns** displays a performance chart so that you can identify when disk, volume, internal volume, and backend volumes spikes or troughs occur during a single day and see typical performance data for different time periods.
- Topology

The analysis of the performance data might include information about contention. This information is displayed on specifically identified tabs, such as the Contending Volumes version of the Volumes tab.

If you analyze performance on a storage virtualizer, the Volumes and Internal Volumes tabs show the performance of the virtual volumes and internal volumes created by the virtualizer. The Disks tab shows the performance of the virtual disks used by the virtualizer. The Backend Volumes tab shows the performance of all the backend volumes used by the virtualizer.

Related concepts

[Performance trend analysis](#) on page 18

[Automatic storage tiering in Inventory and Performance views](#) on page 67

Related tasks

[Searching for performance data](#) on page 16

[Monitoring data store performance](#) on page 24

[Troubleshooting slow client computers on your network](#) on page 43

[Allocating capacity to a new host](#) on page 60

Related references

[Host Summary tab](#) on page 75

Daily Performance Pattern histogram

The Daily Patterns tab displays a histogram showing performance data for the selected disk, volume, end volume, internal volume, or backend volume as a percentage utilization rate. You can change the time period for the performance data sample and refresh the histogram.

Navigation

From the OnCommand Insight Open menu, select **Performance** and any of the performance views. Right-click on a device in the performance view and select the Analyze option. Click the **Daily Patterns** tab to display the performance histogram.

Descriptions

Performance Data You select from these options:

- Disk
- Internal Volume
- Volume
- Backend Volume

Metric You select from the utilization, IOPS, Throughput, and Accessed options in the pull-down menu.

Data Store Summary tab

The Data Store Summary tab on the Analyze dialog box provides performance troubleshooting data, enables you to see the basic information for the element and the resources it is using, and links to additional information for the storage resources.

Troubleshooting information

The left side of this summary lists:

- Name of the data store
- Capacity
- VMDKs capacity
- Cluster Hosts:
 - CPU average utilization
 - Memory utilization
- Cluster FC Network:
 - Port utilization
 - Port errors
- Storage resource used by the data store
 - Links to individual storage resources for more information
 - Capacity of the storage resource
 - Response Time
 - IOPS
 - Throughput

Numbers shown in red on the Data Store Summary tab indicate values that are beyond the policy threshold.

Network information

The right side of this summary lists:

- Virtual Center IP address with a link to more information
- Data sources

Possible error messages

Two error messages might display on this tab:

- Data is not being collected indicates that the data sources have stopped collecting data.
- Data last collected (time) indicates that the displayed data is stale.

Related tasks

[Monitoring data store performance](#) on page 24

Disks tab

You use the Disks tab in the Analyze dialog box to see all of the arrays and the corresponding disks that the hosts are using. You can identify any hosts that are competing for array resources. This information helps you understand potential contention issues that might arise. You can change the analysis period and refresh the data.

Navigation

- To display the Disks tab in the Analyze dialog box, right-click one or more items in any of these views:
 - Storage Arrays
 - Application Performance
 - Switch Port Performance
 - Storage Performance
- Select Analyze.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	Name of the disk in the array.
Storage	Name of the storage array.
<i>icon (Is virtual?)</i>	A "V" icon indicates that the device is a virtualized volume.
Hosts	Name of the hosts using the disk.
Disk Size (GB)	The physical disk capacity, in gigabytes. This differs from usable capacity when technologies such as RAID are used, where some of the raw capacity is used for protection purposes.

Disk Group	Name of the set of disks that share a common configuration. A disk group contains volumes that must use disks within that disk group.
Disk Speed (RPM)	How fast the disk spins, in RPM. (This value is not applicable to SSD disks.)
Utilization	The portion of the total percentage of pre-cache Service Time used for Read and Write requests out of the selected time range.
Utilization (R&W)	The percentage of pre-cache Service Time used for Read or Write requests out of the selected time range.
Position	Where the disk is located (for example, as in Symmetrix: DirectorID-075, Slot 38).
Role	Role of the disk, as obtained from the vendor. Values include data, faulty, parity, double parity, mixed, spare, unallocated, or other. For spare disks, "spare" is displayed.
Status	Status of the disk. The following values could appear: Normal, copying, Failed, Partner, Pending, Questionable, Reconstructing, Zeroing, Other.
Model	Name of disk model.
Vendor	Name of disk vendor.
Type	Type of disk. One of the following values appears: ATA, FATA, FC, SAS, SATA, SATA2, SATA3, SSD, Other.
Serial Number	Serial number of the disk, if the data source supports reporting it.
Top Utilization	The maximum percentage of the total pre-cache Service Time that can be used for requests out of the selected time range.
IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).
IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum number of I/O service requests that can pass through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Throughput	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Throughput (R&W)	Rate at which Read or Write data is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Throughput	Maximum rate at which data can be received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Accessed	"Yes" indicates that the resource was accessed in the selected time range.
Application	Name of application.

Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.

Host Summary tab

The Host Summary tab on the Analyze dialog box provides performance troubleshooting data for the hosts, enables you to see the basic information for the element and the resources it is using, and links to additional information for the storage resources.

Troubleshooting information

The left side of this summary lists:

- Name of the host
- IP addresses
- Storage Resources:
 - Total capacity
 - Links to individual storage resources for more information
 - Capacity of the storage resource
 - Response time
 - IOPS
 - Throughput
- FC ports used by the cluster network:
 - Balance index
 - Name of the FC port
 - Utilization of the FC port
 - Transfer Rate
 - Connected to
 - Errors
- Recent Violation Events

Numbers shown in red on the Host Summary tab indicate values that are beyond the policy threshold.

Location and application information

The right side of this summary lists:

- Data Center name
- Applications running on the host
- Business Entities supported on the host
- Auto Resolution indicator

Possible error messages

Two error messages might display on this tab:

- `Data is not being collected` indicates that the data sources have stopped collecting data.
- `Data last collected (time)` indicates that the displayed data is stale.

Related tasks

[Searching for performance data](#) on page 16

[Troubleshooting slow client computers on your network](#) on page 43

[Troubleshooting poor application performance](#) on page 44

[Analyzing port balance violations](#) on page 53

Related references

[Host Performance view](#) on page 93

Internal Volume Summary tab

The Internal Volume Summary tab on the Analyze dialog box provides performance troubleshooting data, enables you to see the basic information for the element and the resources it is using, and links to additional information for the internal volumes.

Troubleshooting information

The left side of this summary lists:

- Name of the internal volume
- Name of the storage on which the storage pool resides
- Used capacity
- Snapshot capacity
- Deduplication
- Top response time
- Top IOPS
- Top Throughput
- Storage Pools used by the internal volume:
 - Link to information about a storage pool
 - Response Time of the storage pool
 - IOPS
 - Throughput

Numbers shown in red on the Internal Volume Summary tab indicate values that are beyond the performance policy threshold.

Internal volume and application information

The right side of this summary lists:

- Type of volume
- Network technology
- Number of shares
- Clone source link
- Tier
- Applications running on the internal volume
- Business Entities supported on the internal volume
- Data sources

Possible error messages

Two error messages might display on this tab:

- `Data is not being collected` indicates that the data sources have stopped collecting data.
- `Data last collected (time)` indicates that the displayed data is stale.

Storage Pool Summary tab

The Storage Pool Summary tab on the Analyze dialog box provides performance troubleshooting data, enables you to see the basic information for the element and the resources it is using, and links to additional information for the storage pools.

Troubleshooting information

The left side of this summary lists:

- Name of the storage pool
- Storage pool vendor description
- Used capacity
- Top Response Time
- Top IOPS
- Top Throughput
- Resources provisioned for the storage pool with the number of resources and the Provisioned capacity (GB of total and percentage):
 - Link information about a resource
 - Response Time for the resource
 - IOPS of the resource
 - Throughput of the resource
- Disks used by the storage pool:
 - Disk size
 - Disk type

- Disk speed
- Average utilization
- List of the individual disks:
 - Name of the disk
 - Utilization of the disk
 - IOPS
 - Throughput

Numbers shown in red on the Storage Pool Summary tab indicate values that are beyond the performance policy threshold.

Storage Pool information

The right side of this summary lists:

- Type of storage pool
- Redundancy
- Tier
- Data center
- Data sources

Possible error messages

Two error messages might display on this tab:

- `Data is not being collected` indicates that the data sources have stopped collecting data.
- `Data last collected (time)` indicates that the displayed data is stale.

Virtual Machine Summary tab

The Virtual Machine Summary tab on the Analyze dialog box provides performance troubleshooting data, enables you to see the basic information for the element and the resources it is using, and links to additional information for the virtual machines.

Troubleshooting information

The left side of this summary lists:

- Name of the virtual machine
- CPU Utilization
- Memory utilization
- Hypervisor:
 - Link to information about the Hypervisor
 - Hypervisor CPU Count
 - CPU Utilization
 - Memory Utilization

- VM Count
- VMDKs:
 - Total VMDK Capacity
 - Name of the VMDK
 - Capacity
 - Used Capacity
 - Latency
 - Link to the VMDK storage resource
- Datastores used by the virtual machine:
 - Link to the data store by Name
 - Capacity of each data store

Numbers shown in red on the Virtual Machine Summary tab indicate values that are beyond the performance policy threshold.

Network and application information

The right side of this summary lists:

- Applications running on the virtual machine
- Business Entities supported on the virtual machine
- Data sources

Possible error messages

Two error messages might display on this tab:

- `Data is not being collected` indicates that the data sources have stopped collecting data.
- `Data last collected (time)` indicates that the displayed data is stale.

Related tasks

[Determining if a VM is affecting host performance](#) on page 41

Volume Summary tab

The Volume Summary tab on the Analyze dialog box provides performance troubleshooting data, enables you to see the basic information for the element and the resources it is using, and links to additional information for the volumes.

Troubleshooting information

The left side of this summary lists:

- Name of the volume
- Storage type

- Capacity
- Top Response Time
- Top IOPS
- Top Throughput
- Clients using the volume:
 - Link to client information
 - Response Time for the client
 - IOPS of clients
 - Throughput of the client
- FC storage ports mapped to the volume:
 - Name of the port
 - Utilization
 - Transfer Rate of the port
 - Connected to
 - Errors
- Internal Volume used by the volume:
 - Link to information about an internal volume
 - Response Time of the internal volume
 - IOPS
 - Throughput
- Storage Pools used by the volume:
 - Link to information about a storage pool
 - Response Time of the storage pool
 - IOPS
 - Throughput

Numbers shown in red on the Volume Summary tab indicate values that are beyond the performance policy threshold.

Volume and application information

The right side of this summary lists:

- Type of volume
- Tier
- Applications running on the volume
- Business Entities supported on the volume
- Data sources

Possible error messages

Two error messages might display on this tab:

- `Data is not being collected` indicates that the data sources have stopped collecting data.
- `Data last collected (time)` indicates that the displayed data is stale.

Related tasks

[Troubleshooting poor application performance](#) on page 44

[Searching for performance data](#) on page 16

Application Performance view

In the Application Performance view, you can display performance data for applications to determine whether one application has affected the performance of another application. You might display information in this view by the business entities you have defined and associated with applications.

Navigation

From the OnCommand Insight Open menu, select **Performance > Application Performance**.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Application	Name of application.
Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
Volume Response Time	The time it takes from the moment a request for information arrives at the storage device to the time when the storage device begins to send the information back in response.
Volume Response Time (R&W)	The time it takes from the moment a request for information arrives at the storage device to the time when the storage device begins to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Top Volume Response Time	The longest time it takes from the moment a request for information arrives at the storage device until the time when the storage device begins to send the information back in response.
Volume IOPS	The number of Read or Write I/O service requests.

Volume IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top Volume IOPS	The maximum number of Read or Write I/O service requests.
Volume Throughput	Rate at which data was transferred to/from the volume in response to I/O service requests during the time period of the selected item.
Volume Throughput (R&W)	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Volume Throughput	The maximum rate that data can be transmitted.
Volume Cache Hit Ratio (R&W)	Percentage of Read/Write requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better the performance. This column is empty for storage arrays that do not collect cache hit information.
Volume Partial (R&W)	Total number of times that a read or write crosses a stripe boundary on the volume. Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size. Indication of improper alignment of a volume (or a NetApp LUN).
Write Pending	The number of Write I/O service requests that are pending.
Volume Accessed	“Yes” indicates that the resource was accessed in the selected time range.
Internal Volume Response Time	The time it takes from the moment a request for information arrives at the storage device to the time when the storage device begins to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. The OnCommand Insight response time value is measured in milliseconds from within the storage array.
Internal Volume Response Time (R&W)	The time it takes from the moment a request for information arrives at the storage device to the time when the storage device begins to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Top Internal Volume Response Time	The maximum time it takes from the moment a request for information arrives at the storage device until the time when the storage device begins to send the information back in response.
Internal Volume IOPS	The total number of Read/Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).

Internal Volume IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top Internal Volume IOPS	The maximum number of Read or Write I/O service requests.
Internal Volume Throughput (R&W)	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Internal Volume Throughput	The maximum rate that data can be transmitted in a fixed amount of time.
Internal Volume Accessed	“Yes” indicates that the resource was accessed in the selected time range.

Options

From the Application Performance view, right-click to show a pop-up menu containing the following options.

Analyze	Available only with the Perform license. Allows you to investigate the performance of the selected resources affected by the violation. For example, you can determine contention issues, availability issues, and array performance.
Analyze Contention	Available only with the Perform license. Allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are in contention for performance on the disk, a factor that can cause poor performance.
Analyze Storage Pools	Available only with the Assure license. Allows you to select a specific storage pool and assess its status related to the thin-provisioning policies. You can use this dialog box, instead of the Violations Browser, to see the current thin provisioning violations and how close the storage pool is to reaching the policy limits.

Backend Volume Performance view

Using the Backend Volume Performance view, you can see all the volumes residing on the backend storage arrays connected to the selected virtualizer instead of the virtual volumes. This data helps you better use capacity and assists in proactive analysis of how the storage resources are being used. You can display information in this view by the business entities you have defined and associated with applications.

Navigation

From the OnCommand Insight Open menu, select **Performance > Application Performance**. Click the **Backend Volume Performance** icon.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Volume	Name of volume.
Label	An alternate name or alias assigned to a volume by the storage administrator.
UUID	Universally unique identifier for the object. In this case, it is generated by and retrieved from the storage array itself.
Storage	Name of the storage array.
Storage Pool	The name of the storage pool on which the backend volume resides.
Auto Tiering	A checkmark indicates that the selected storage pool is using the automatic storage tiering technology (for example, FAST VP).
Thin Provisioned	A checkmark indicates that the volume is leveraging thin provisioning.
vFiler	The name of the vFiler unit. A vFiler unit is an isolated software container that behaves exactly like a physical storage array. A vFiler unit shares the physical resources of the array, but abstracts the client access from the physical array into virtual arrays.
Internal Volume	Name of the internal volume that the volume uses.
Qtree	Name of the qtree on this volume.
<i>icon (array virtualization type)</i>	Indicates type of virtualization. Showing on a virtual volume, a "V" icon indicates that the device is a virtualized volume and a "B" icon indicates that the device is a backend volume.
Virtualizer	For backend volumes in array virtualization. Displays the name of the front end virtualizer that is using this volume.
Virtual Storage Pool	For backend volumes in array virtualization. The name of the storage pool on the front end virtualizer that is using this volume.
Datastore	The name of data store residing on this volume.
Application	Application associated with this volume.
Application Priority	The importance of this application within your organization that was set when defining the application.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
Hosts	Hosts making use of the volume.
Capacity (GB)	Size of the volume that is accessible to host applications, in gigabytes.

Raw Capacity (GB)	Physical disk capacity of the volume, in gigabytes. This differs from usable capacity when technologies such as RAID-5 are used, where some of the raw capacity is used for protection purposes.
Consumed Capacity (GB)	The amount of capacity that the volume consumes from underlying storage (for example, internal volume and storage pool). For non-thin provisioned value, that is the same as the volume's capacity. For thin provisioned volumes, it is the amount of capacity used to store the volume's contents. Its value is also affected by Snapshot copies, deduplication, and other storage technologies. If the volume does not have usage information available, this appears blank.
Mapped Ports	The number of storage ports through which this volume is accessible.
Redundancy	Level of mirroring defined for the device based on the storage technology, for example, RAID-DP, underlying the device. This is taken from the device itself. For an explanation of values, see the device documentation.
Response Time	The time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response.
Response Time (R&W)	The time it takes for a read or write request to arrive at the storage device and to respond to the request.
Top Response Time	The maximum length of time in milliseconds that it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response.
IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).
IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum number of I/O service requests that can pass through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Throughput	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Throughput (R&W)	Rate at which Read or Write data is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Throughput	Maximum rate at which data can be received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Cache Hit Ratio	Percentage of requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better the performance. This column is empty for storage arrays that do not collect cache hit information.

Cache Hit Ratio (R&W)	Percentage of Read/Write requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better the performance. This column is empty for storage arrays that do not collect cache hit information.
Partial (R&W)	Total number of times that a read or write crosses a stripe boundary on any disk module in RAID 5, RAID 1/0 or RAID 0 LUN. Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size and is an indication of improper alignment of a volume (or a NetApp LUN). For CLARiiON only, this value is the number of stripe crossings divided by the total number of IOPS.
Write Pending	The number of Write I/O service requests that are pending.
Accessed	"Yes" indicates that the resource was accessed in the selected time range.
<i>annotations</i>	Annotations associated with each volume.

Datastore Performance view

This view correlates data store performance with storage and switch port performance. The data store performance metrics are aggregates of the VMDK performance. You can see the sum of IOPS and throughput, the maximum latency values, and the average Hypervisor CPU and memory utilization.

Navigation

From the OnCommand Insight Open menu, select **Performance > Datastore Performance**.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	The data store name that represents the virtual machines.
Virtual Center IP	The IP address of the Virtual Center Host for the data store.
VM Count	The number of virtual machines whose files are contained in this data store.
Hypervisor Count	The number of Hypervisor hosts that use this data store for their virtual machines.
FC Ports	The number of fibre channel ports in the physical storage paths that this data store logically represents.
Capacity (GB)	Usable capacity or configured size of the data store, in gigabytes.

Provisioned Capacity (GB)	The amount of total capacity that has been set aside for potential use, based on the virtual machines using this data store. Includes space set aside for virtual machine files of all types.
Used Capacity (GB)	The amount of capacity holding actual data in the data store. Includes usage based on all file types.
Unused Capacity (GB)	The available capacity of the data store, in gigabytes.
VMDKs Capacity (GB)	The virtual machine usable capacity, in gigabytes.
Over-committed Capacity (GB)	The amount of capacity that has been overcommitted from this resource. When thin provisioning is in use, the total size can exceed the total size of the capacity committed. If there is no over commitment on the data store, the value is 0.
Commit Ratio	The ratio of the sum of the capacity of all virtual disks allocated on a data store to the capacity of the data store. If thin provisioning is in place, the rate can be greater than 100% meaning it is over committed.
Has FC Port Errors	Indicates if there were any port errors over the given time span.
FC Port Utilization	The average utilization percent for the FC ports belonging to the data store's physical paths over the given time span.
Storage	Storage arrays used by this data store.
Resource Name	The volumes or internal volumes in the path.
Resource Capacity (GB)	The total capacity, in gigabytes, of the volumes or internal volumes in this data store's paths.
Resource Used Capacity (GB)	The total used capacity, in gigabytes, of the volumes or internal volumes in this data store's paths.
Resource Technology	The SAN (FC and iSCSI) or NAS (NFS and CIFS) protocols that the device supports.
Deduplication Savings	The known amount of storage savings through deduplication, a process that detects blocks with identical content and replaces subsequent identical blocks with a reference to a single copy of the block.
Storage Pools	Storage pools used by this data store.
HV Memory Utilization	The average memory utilization percent of this data store's Hypervisors for the given time span.
HV CPU Utilization	The average memory utilization percent of this data store's Hypervisors for the given time span.

VMDK IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
VMDK IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).
VMDK Top IOPS	The highest of the IOPS reported by the measured devices.
VMDK Throughput (R&W)	Rate at which Read or Write data is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
VMDK Throughput	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
VMDK Top Throughput	The highest of the throughputs reported by the measured devices.
VMDK Latency (R&W)	The response time for a Read or Write from the virtual machines.
VMDK Latency	The average response time from the virtual machines carved from a data store.
VMDK Top Latency	The highest response time from the virtual machines carved from a data store.
Storage IOPS	The average IOPS of the data store's storages for the given time span.
Storage IOPS (R&W)	The average read or write IOPS for the data store.
Top Storage IOPS	The maximum observed IOPS (peak) for all data stores within the hour.
Storage Response Time	The average response time of the data store's storage for the given time span. This is the time it takes from the moment a request for information is sent until the information starts arriving.
Storage Response Time (R&W)	The maximum read or write response time for this data store throughout the hour.
Top Storage Response Time	The maximum observed response time for this data store within the hour.
Partial R/W	Total number of times that a read or write crosses a stripe boundary on any disk module in RAID 5, RAID 1/0 or RAID 0 LUN. Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size and is an indication of improper alignment of a volume (or a NetApp LUN). For CLARiiON only, this value is the number of stripe crossings divided by the total number of IOPS.

Storage Throughput	The average throughput of the data store's storage for the given time span.
Storage Throughput (R&W)	The average read or write throughput for the given time.
Top Storage Throughput	The maximum observed (peak) throughput for the data store's storage for a given time.

Options

The following options are available from the right-click menu:

Analyze	Available only with the Perform license. Allows you to investigate the performance of the selected resources affected by the violation. For example, you can determine contention issues, availability issues, and array performance. The Data Store Summary tab provides information that might be needed for troubleshooting.
Analyze Contention	Available only with the Perform license. Allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are in contention for performance on the disk, a factor that can cause poor performance.
Analyze Storage Pools	Available only with the Assure license. Allows you to select a specific storage pool and assess its status related to the thin provisioning policies. You can use this dialog box, instead of the Violations Browser, to see the current thin provisioning violations and how close the storage pool is to reaching the policy limits.
Modify Policy	Changes the policy that governs alerts set on this resource. This option requires the Assure license.

Related tasks

[Monitoring data store performance](#) on page 24

Disk Performance view

You use the Disk Performance view to see all the arrays and corresponding disks that the hosts are using and metrics that can help you analyze performance. You can identify any hosts that are competing for array resources. This view can be vital because it helps you see potential contention issues before they arise.

Navigation

From the OnCommand Insight Open menu, select **Performance > Host Performance** or **Datastore Performance**. Click the **Disk Performance** icon.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	Name of the disk in the array.
Storage	Name of the storage array.
<i>icon (Is virtual?)</i>	A "V" icon in this column indicates that the device is a virtualized volume.
Hosts	Name of the hosts using the disk.
Disk Size (GB)	The physical disk capacity, in gigabytes. This differs from usable capacity when technologies such as RAID are used, where some of the raw capacity is used for protection purposes.
Disk Group	Name of the set of disks that share a common configuration. A disk group contains volumes that must use disks within that disk group.
Disk Speed (RPM)	How fast the disk spins, in revolutions per minute. (This value is not applicable to SSD disks.)
Utilization	The portion of the total percentage of pre-cache Service Time used for Read and Write requests out of the selected time range.
Utilization (R&W)	The percentage of pre-cache Service Time used for Read or Write requests out of the selected time range.
Position	Where the disk is located (for example, as in Symmetrix: DirectorID-075, Slot 38).
Role	Role of the disk, as obtained from the vendor. Values include data, faulty, parity, double parity, mixed, spare, unallocated, or other. For spare disks, "spare" appears.
Status	Status of the disk. The following values could appear: Normal, Copying, Failed, Partner, Pending, Questionable, Reconstructing, Zeroing, Other.
Vendor	Name of disk vendor.
Type	Type of disk. One of the following values appears: ATA, FATA, FC, SAS, SATA, SATA2, SATA3, SSD, Other.
Model	Name of disk model.
Serial Number	Serial number of the disk, if the data source supports reporting it.
Top Utilization	The maximum percentage of the total pre-cache Service Time that can be used for Read and Write requests out of the selected time range.
IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).

IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum number of I/O service requests that can pass through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Throughput	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Throughput (R&W)	Rate at which Read or Write data is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Throughput	Maximum rate at which data can be received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Accessed	"Yes" indicates that the resource was accessed in the selected time range.
Application	Name of application.
Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.

Options

The following options are available from the right-click menu:

Show Storage Pools	The Storage Pools by Disk dialog box lists the capacity information, disk size, disk group, and disk speed for each disk.
Analyze	Available only with the Perform license. Allows you to investigate the performance of the selected resources affected by the violation. For example, you can determine contention issues, availability issues, and array performance.
Analyze Contention	Available only with the Perform license. Allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are contending for performance on the disk, a factor that can cause poor performance.
Analyze Storage Pools	Available only with the Assure license. Allows you to select a specific storage pool and assess its status related to the thin provisioning policies. You can use this dialog box instead of the Violations Browser to see the current thin provisioning violations and how close the storage pool is to reaching the policy limits.

Related concepts

[Automatic storage tiering in Inventory and Performance views](#) on page 67

Related tasks

[Allocating capacity to a new host](#) on page 60

[Allocating capacity to an existing host](#) on page 62

Histogram view

The OnCommand Insight histograms use a bar chart to show the relative amount of time the selected ports spent in the given utilization levels. For the daily traffic, the histogram identifies when there are traffic pattern spikes or troughs during a single day and identifies the average hourly values for the analysis period. The performance histogram shows the percentage of available bandwidth used for transmitted (Tx) and received (Rx) data. The Analyze dialog box supplies a specialized histogram with multiple data type selections.

Navigation

- From the OnCommand Insight Open menu, select **Performance > Switch Port Performance**. Right-click on a device in the performance view and select either the **Show Daily Traffic Pattern** or **Show Performance Histogram** option.
- Open the Analyze dialog box and click the **Daily Patterns** tab to display the Daily Performance Pattern histogram. This histogram supplies selections for performance data and metrics to focus the data displayed in the bar chart.

Descriptions

You can sort and filter the data in the histograms as in other OnCommand Insight views.

The histogram for the **Switch Port Performance** data includes Fabric, Switch, and Switch Port columns, always reflecting a port on a switch. The Device Type, Device Name, and Device Port show the information about the port on the other end of the path. You can select a Fixed or Variable scale for the display. The **Fixed** selection always begins with zero percent. The **Variable** selection allows you to slide an indicator to begin the data with a higher percentage.

You can select multiple ports and see the collective impact of traffic on the switch or check individual ports for periods of high or low traffic. The first column of the table allows the selection of individual ports. The histogram displays all of the visible, selected rows in the table. Selections in the table at the top control what is displayed in the chart below.

The Daily Performance Pattern histogram displayed from the Analyze dialog box allows you to select different time periods available for the dialog box, performance metrics, and these performance data options:

- Disk
- Volume
- Internal Volume
- Backend Volume

Host Performance view

Use this view to identify multipath validation and host balancing. The performance values aggregate performance statistics for each host. The Internal Volume metric columns show the aggregate performance measurements of all internal volumes that the host accesses. The Volume metric columns show the aggregate performance measurements of all volumes the host accesses.

Navigation

From the OnCommand Insight Open menu, select **Performance > Host Performance**.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
<i>registered icon</i>	Icon indicating that the host referenced by the policy is registered to the current user. To view only those policies for your registered hosts, filter by this icon.
Name	Name of host on the network.
<i>icon (Are there violations?)</i>	Icon that indicates a policy violation.
Application	Names of the applications associated with the host.
Application Priority	The importance of this application within your organization, as set when defining the application.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
V-Cluster	Name of a cluster of virtualization hosts that share access to the same SAN volumes or NAS share. For a standalone host, this is blank.
CPU Utilization	Amount of actively used CPU, as a percentage of total available (over all virtual CPUs).
Memory Utilization	Threshold for the memory used by the host.
Disk IOPS (R)	Total number of read I/O service requests to the disk passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Disk Latency (R)	The sum of Disk Read Latency for the sampling period. Average amount of time taken during the interval to process SCSI read commands issued from the Guest OS to the virtual machine.

Disk Throughput (R)	Portion of data that the disk read in response to I/O service requests to the disk (measured in MB per sec).
Disk IOPS	The sum of Disk Read IOPS and Disk Write IOPS.
Top Disk IOPS	The maximum (over all disks) for total IOPS.
Disk Latency	The sum of Disk Read Latency and Disk Write Latency for the sampling period.
Top Disk Latency	The maximum disk latency that the disk discovered for the sampling period.
Disk Throughput	Portion of data that the disk sent/received in response to I/O service requests to the disk (measured in megabytes per sec).
Top Disk Throughput	The maximum amount of data that the disk can transmit in a fixed amount of time in response to I/O service requests to the disk (measured in megabytes per sec).
Disk IOPS (W)	Total number of write I/O service requests to the disk passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Disk Latency (W)	The sum of Disk Write Latency for the sampling period. Average amount of time taken during the interval to process SCSI write commands issued from the Guest OS to the virtual machine.
Disk Throughput (W)	Portion of data that the disk writes in response to I/O service requests to the disk (measured in megabytes per sec).
IP Throughput (Receive)	Average rate at which disk IP data was received in megabytes.
IP Throughput	Aggregate rate at which disk IP data was transmitted and received in megabytes.
IP Throughput (Transmit)	Average rate at which disk IP data was transmitted in megabytes.
Disk Accessed	Indication if at least one disk was accessed (for read or write).
IP Accessed	Indication if Guest OS received or transmitted any IP data.
FC Balance Index	The standard deviation from the traffic average for all of the Fibre Channel ports; a high value (above 50) signifies a ratio problem between the host ports, meaning that some ports are experiencing a large amount of traffic and others very low, indicating a potential problem.
Volume Response Time (R&W)	The time in milliseconds that it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible.

	OnCommand Insight measures response times from within the storage array.
Top Volume Response Time	The maximum amount of time it can take from the moment a request for information arrives at the storage device until the time when the storage devices begin to send the information back in response.
Volume IOPS (R&W)	Measures the total number of Read/Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top Volume IOPS	The maximum number of Read or Write I/O service requests.
Volume Throughput (R&W)	Amount of Read or Write data that is being received in a fixed amount of time in response to I/O service requests (measured in megabytes per sec).
Top Volume Throughput	The maximum rate that data can be transmitted.
Volume Cache Hit Ratio (R&W)	Percentage of Read/Write requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better the performance.
Volume Partial R/W	Total number of times that a read or write crosses a stripe boundary on the volume. Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size, while a high percentage indicates improper alignment of a volume (or a NetApp LUN).
Write Pending	The number of Write I/O service requests that are pending.
Volume Accessed	“Yes” indicates that the resource was accessed in the selected time range.
Internal Volume Response Time (R&W)	The time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Top Internal Volume Response Time	The maximum time it takes from the moment a request for information arrives at the storage device until the time when the storage devices begin to send the information back in response.
Internal Volume IOPS (R&W)	The total number of Read/Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top Internal Volume IOPS	The maximum number of Read or Write I/O service requests.
Internal Volume Throughput (R&W)	Amount of Read or Write data that is being received in a fixed amount of time in response to I/O service requests (measured in megabytes per sec).

Top Internal Volume Throughput The maximum rate that data can be transmitted in a fixed amount of time.

Internal Volume Accessed “Yes” indicates that the resource was accessed in the selected time range.

annotations User-defined terminology associated with the devices.

Options

The following options are available from the right-click menu:

- Analyze** Available only with the Perform license. Allows you to investigate the performance of the selected resources affected by the violation. For example, you can determine contention issues, availability issues, and array performance. The Host Summary tab provides information that might be needed for troubleshooting.
- Analyze Contention** Available only with the Perform license. Allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are in contention for performance on the disk, a factor that can cause poor performance.
- Analyze Storage Pools** Available only with the Assure license. Allows you to select a specific storage pool and assess its status related to the thin provisioning policies. You can use this dialog box, instead of the Violations Browser, to see the current thin provisioning violations and how close the storage pool is to reaching the policy limits.
- Modify Policy** Changes the policy that governs alerts set on this resource. This option requires the Assure license.

Related tasks

[Allocating capacity to an existing host](#) on page 62

[Troubleshooting slow client computers on your network](#) on page 43

[Troubleshooting poor application performance](#) on page 44

Related references

[Host Summary tab](#) on page 75

Internal Volume Performance view

You can view the internal volume relevant to the object selected and determine the internal volumes that are competing for storage resources. Performance values reflect an aggregation of the load on all of its shares and volumes residing on the internal volume.

Navigation

From the OnCommand Insight Open menu, select **Performance > Host Performance**. Click the **Internal Volume Performance** icon.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	Name of internal volume.
Storage	Name of storage array on which the internal volume resides.
vFiler	The name of the vFiler unit. A vFiler unit is an isolated software container that behaves exactly like a physical storage array. A vFiler unit shares the physical resources of the array, but abstracts the client access from the physical array into virtual arrays.
Application	Applications associated with the internal volume.
Application Priority	The importance of this application within your organization that was set when defining the application.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
Hosts	Hosts associated with the internal volume.
Capacity (GB)	Usable capacity of the internal volume, in gigabytes.
Raw Capacity (GB)	The physical disk capacity of the internal volume, in gigabytes. This differs from usable capacity when technologies such as RAID are used, where some of the raw capacity is used for protection purposes.
Used Capacity (GB)	The amount of capacity holding actual data in the internal volume. Includes usage based on all file types.
Consumed Capacity (GB)	The amount of capacity that the volume consumes from underlying storage (for example, internal volume and storage pool). For non-thin provisioned volumes, this value is the same as the volume's capacity. For thin provisioned volumes, it is the amount of capacity used to store the volume's

contents. Its value is also affected by Snapshot copies, deduplication, and other storage technologies. If the volume does not have usage information available, this appears blank.

Storage Pool	The name of the storage pool on which the internal volume resides.
Datastore	The name of data store residing on this volume.
Type	Type of internal volume, for example, FlexClone or FlexVol.
Thin Provisioned	A check mark indicates that the resource is thin provisioned.
Performance Policy	The level of performance threshold (for example, Global) set on this resource. For example, maximum IOPS, response time, or throughput thresholds might be set for specific resources and not use the global thresholds.
Space Guarantee	A vendor-specific indication of how the internal volume's space is allocated if thin provisioning is used.
Deduplication Savings	Percentage value of the rate of deduplication in effect for the internal volume.
Clone Source	The name of the internal volume that this internal volume cloned. This is the source of the cloned relationship.
Clone Shared Capacity (GB)	When the Internal Volume is a clone of another Internal Volume, this value shows the amount of capacity that is shared between the Source and Target clones (If no changes have been made to the source and target Internal Volumes since the clone operation occurred, all capacity will be shared. Modifications made to either of the Internal Volumes will decrease the shared capacity).
Status	Information about whether the internal volume is online, offline, or OK.
Snapshot Reserve (GB)	The capacity that was reserved for Snapshot copies.
Snapshot Used (GB)	The capacity remaining after some capacity that was reserved for Snapshot copies was actually used.
Snapshot Used (%)	The percent of capacity remaining after some capacity that was reserved for Snapshot copies was actually used.
Snapshot Overflow (GB)	The capacity that was used for Snapshots, but exceeded what was reserved for Snapshot copies.
Snapshot Count	The number of Snapshot copies that are stored for this internal volume.
Last Snapshot	The time when the last Snapshot copy operation occurred on this internal volume.

Disk Types	The types of physical disks (for example, Fibre Channel or ATA) on which the internal volume is based.
Disk Size (GB)	The size of the physical disks on which the internal volume is based.
Disk Speed (RPM)	The speed (RPM) of the physical disks on which the internal volume is based.
Response Time	The time that it took to access the disk. These values are based on fixed speed specifications of the disks.
Response Time (R&W)	The time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Top Response Time	The maximum length of time it took for the internal volume to begin to send the response.
IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).
IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum number of I/O service requests generated by the internal volume over the selected length of time.
Throughput (R&W)	Rate at which data is read or written in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Throughput	Maximum rate at which Read or Write data can be received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Accessed	“Yes” indicates that the resource was accessed in the selected time range.
<i>annotations</i>	User-defined terminology associated with the internal volumes including Service Level, Note, and Tier.

Internal Volume Usage of Disk view

To examine additional information for each internal volume in the Disk Utilization Violations view, open the Internal Volume Usage of Disk view.

Navigation

From the OnCommand Insight Open menu, select **Assurance > Disk Utilization Violations**. At the bottom of the view, click the **Internal Volume Usage of Disk** icon.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	Name of internal volume.
Hosts	Hosts associated with the internal volume.
Application	Applications associated with this internal volume.
Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
Internal Volume IOPS	Measures the total number of I/O service requests on the internal volume during the time of the selected disk utilization violation (measured in I/O per sec).
Disk IOPS	Measures the total number of I/O service requests on the disk for the volume during the time of the selected disk utilization violation (measured in I/O per sec).
Disk IOPS Distribution	Portion of the disk IOPS that the volume contributed toward the total disk IOPS during the time period of the selected disk utilization violation.
Internal Volume Throughput	Rate at which data was transferred to/from the internal volume in response to I/O service requests during the time period of the selected disk utilization violation (measured in MB per sec).
Disk Throughput	Rate at which data was transferred to/from the disk in response to I/O service requests on the volume during the time period of the selected disk utilization violation (measured in MB per sec).

Disk Throughput Distribution	Portion of the disk throughput that the volume contributed toward the total disk throughput during the time period of the selected disk utilization violation.
Internal Volume Response Time	The time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Internal Volume Top Response Time	The maximum length of time it takes to begin sending back a response.

Performance Chart

You can view performance trends over time to determine the cause of a violation in the Performance Chart. It also allows you to examine the various performance metrics over a specified period and narrow your search using the detailed information displayed when you position the mouse pointer over points in the chart. You can display the Analyze dialog box for those points in the Performance Chart.

Navigation

- From the Host Performance, Virtual Machine Performance, Datastore Performance, Switch Port Performance, or Storage Performance view, click the **Performance Chart** icon.
- In the Analyze dialog box, select one or more items on the Disks, Volumes, Internal Volumes, Backend Volumes, Hosts, Datastores, Virtual Machines, VMDKs, or Switch Ports tab and click the **Performance Chart** icon.

Descriptions

If automatic storage tiering would display incorrect data, no data appears on the chart. For details, click the **Auto tiering affects data** link.

Show Change Events	Select this option to display letters as change event markers on the chart. Clicking on the markers highlights the change event in the list. The change events on the list give a brief description of what the change involved and at what time the change was detected by OnCommand Insight. Clicking on the description link opens the Changes detail view that gives the full details of the changes that occurred at that time point.
Pan	Click the hand icon and move the mouse pointer over the performance chart to display details for any point in the graph and right-click to open the Analyze dialog box for that point.

Legend	Displays the metrics displayed on the chart that show performance thresholds and utilization types.
Time Range	<p>Changes the time period for the chart. You can select the following time ranges: Live Sample, Today, Last Hour, Last 24h, Last Week, All, and Custom.</p> <ul style="list-style-type: none"> • With the Custom range, you can select any time range up to the current time. • With the Live Sample, OnCommand Insight first retrieves data for the last hour, then asks the server for changes from the last update time to the current time based on the following, as long as the chart is visible. <ul style="list-style-type: none"> • VM data: 5 to 15 minutes • Switch performance data: 5 to 15 minutes • Storage arrays data: 5 to 15 minutes • Configuration changes: 40 minutes
Date and Time Stamp on x axis	The x axis displays the date and time when OnCommand Insight sampled and retrieved the data.
Zoom	The y axis automatically scales to 5x greater than the greatest value in the chart. In addition, you can zoom into a portion of the chart using the Zoom tool. To zoom in, click the Zoom tool and click and drag across a section of the chart. A box defines the boundaries of the detail that you will zoom into. To zoom out, hold down the Shift key and drag across a region. To reset the default zoom, click the Zoom to Fit tool. These settings remain each time you view the Performance Chart.

Move the mouse pointer over the chart points to display additional information about that specific point in the chart. You might also want to right-click on a chart point and display the Analyze information for it. If you position the mouse pointer on an item in the legend, the minimum and maximum points for that part of the chart are shown.

Related concepts

[Performance trend analysis](#) on page 18

Related tasks

[Identifying performance trends](#) on page 30

[Analyzing virtual storage performance](#) on page 50

Performance Distribution view

The Performance Distribution view illustrates traffic distribution on the selected device using a variety of the distribution analysis metrics.

Navigation

From the OnCommand Insight Open menu, select **Performance** and the Host Performance, Virtual Machine Performance, Datastore Performance, or Switch Port Performance view. Select one or more lines in the view. Click the **Performance Distribution** icon.

Descriptions

Each selected device is represented by a different color in the distribution chart. The legend identifying each color is to the right of the distribution chart.

Changing the selected devices in the view changes the distribution display to the newly-selected devices.

Move the mouse pointer over the distribution chart points to display additional information about specific points in the chart.

If automatic storage tiering would display incorrect data, no data appears on the chart. For details, click the **Auto tiering affects data** link.

Metric For the Switch Port Performance data, you can filter the traffic data displayed on the chart by selecting one of the following metrics:

- Traffic
- Loss of Sync
- Loss of Signal
- Tx Traffic
- Rx Traffic

Related concepts

[Performance trend analysis](#) on page 18

[Automatic storage tiering in Inventory and Performance views](#) on page 67

Related tasks

[Viewing virtual machine performance and utilization](#) on page 34

[Exporting performance data as an image or file](#) on page 32

Port Performance Distribution view

The Port Performance Distribution view illustrates traffic distribution on the selected port using distribution analysis metrics. You can select the specific traffic or loss of sync or signal metrics that you want charted.

Navigation

From the OnCommand Insight Open menu, select **Performance > Switch Port Performance**. Select one or more of the port descriptions. Click the **Port Performance Distribution** icon to display this chart for the selected ports.

Descriptions

Select a metric for the distribution analysis from these options:

- Traffic Rate
- Rx Traffic Rate
- Tx Traffic Rate
- BB Credit Errors
- Loss of Sync
- Loss of Signal

Each selected port is represented by a different color in the distribution chart. The legend identifying each port color is to the right of the distribution chart.

Changing the selected ports in the Switch Port Performance view changes the distribution display to the newly selected ports.

Move the mouse pointer over the distribution chart points to display additional information about that specific point in the chart.

Related tasks

[Monitoring Fibre Channel switch performance](#) on page 23

[Selecting an active path for a virtual machine data store](#) on page 39

[Identifying out-of-balance network traffic across host ports](#) on page 55

Settings for Host Virtualization policies

You can set a policy that monitors for active/passive volume multipathing, for identical volume access for hosts in a virtual cluster, and for identical volume LUN for virtual cluster hosts.

Navigation

From the OnCommand Insight Client menu, select **Tools > Settings**. In the left tree of the Settings view, click the **Policies > Host Virtualization** option.

Field descriptions

The Host Virtualization policy is enabled by default on all hosts.

Single storage processor for volume active paths

Alerts the SAN administrator that a volume is being actively accessed via multiple storage processors. For storage arrays that maintain active/passive multipathing, this will result in storage performance degradation (and even service interruption).

Storage models

If you enabled the **Single storage processor for volume active paths** option, select the storage models that are active/passive and that should not allow access through multiple storage processors to the same volume.

Identical volume access for hosts within a virtual cluster

Alerts the SAN administrator that some hosts in a virtual cluster cannot access volumes other members of the group can. Such a configuration may cause these hosts to be unable to run virtual machines that other hosts in the cluster are able to run. (This is also applicable for NAS shares.)

Identical volume LUN for hosts within a virtual cluster

Alerts the storage administrator that a volume is presented to different members of a virtual cluster using different LUNs.

Settings for VM Hosts Thresholds

Use this view to set the range of Hypervisor performance threshold values that show up in green, yellow, or red on the virtual machine Performance views, which display the HV utilization of virtual machine hosts (ESX servers) over a selected period.

Navigation

From the OnCommand Insight Client menu, select **Tools > Settings**. In the left tree of the Settings view, click the **Thresholds > VM Hosts Thresholds** option.

Field descriptions

For each threshold, set a low and high threshold.

- Optimal: Values below the threshold's low value are displayed in green.
- Warning: Values below the threshold's high value are displayed in yellow.
- Error: Other values, such as those above the threshold's high value, are displayed in red.

The settings on this view affect the colors on the following views: Datastore Performance main view, Virtual Machine Performance main view, VMDK view, and Host Performance main view.

Latency (ms)	Average amount of time for a read/write operation from or to the virtual disk before an alert is issued. The default values are Low 20 and High 100.
CPU Utilization (%)	Threshold for the host's CPU MHz being used. The default values are Low 45 and High 75.
Memory Utilization (%)	Threshold for the memory used by the host. The default values are Low 35 and High 65.

Storage Performance view

The Storage Performance view includes information on the storage array’s internal volumes and shows the rounded sum of each internal volume in the array. You use this view to examine the performance details for disks, volumes, and internal volumes.

Navigation

From the OnCommand Insight Open menu, select **Performance > Storage Performance**.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	Name of the storage array.
Technology	The SAN (FC and iSCSI) or NAS (NFS and CIFS) protocols that the device supports.
<i>icon (array virtualization type)</i>	Indicates virtualization and whether the storage acts as a virtualizer device (V) or as backend storage (B).
Capacity (GB)	Total storage array capacity that is accessible to host applications, in gigabytes.

Raw Capacity (GB)	The physical disk capacity, in gigabytes. This differs from usable capacity when technologies such as RAID are used, where some of the raw capacity is used for protection purposes.
Vendor	Name of storage array vendor.
Family	Name of storage array family.
Model	Name of storage array model.
Volume Response Time	The time that it takes to access the volume. These values are based on fixed speed specifications of the volumes.
Volume Response Time (R&W)	The time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Top Volume Response Time	The maximum length of time it can take from the moment a request for information arrives at the storage device until the time when the storage devices begin to send the information back in response.
Volume IOPS	Measures the total number of I/O service requests on the virtual volume during the time period of the selected disk utilization violation (measured in I/O per sec).
Volume IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top Volume IOPS	The maximum number of I/O service requests generated by a single volume (across all storage volumes) over the selected duration.
Volume Throughput	Rate at which data was transferred to/from the volume in response to I/O service requests during the time period of the selected disk utilization violation (measured in MB per sec).
Volume Throughput (R&W)	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Volume Throughput	The maximum rate that data can be transmitted in a fixed amount of time to or from the volume.
Volume Cache Hit Ratio	Percentage of requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better the performance. This column is empty for storage arrays that do not collect cache hit information.

Volume Cache Hit Ratio (R&W)	Percentage of read or write requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better the performance.
Volume Partial R/W	Total number of times that a read or write request crosses a stripe boundary on any disk module in RAID 5, RAID 1/0 or RAID 0 LUN. Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size. Indication of improper alignment of a volume (or a NetApp LUN). For CLARiiON, this is the number of stripe crossings divided by the total number of IOPS.
Write Pending	The number of Write I/O service requests that are pending.
Volume Accessed	“Yes” indicates that the resource was accessed in the selected time range.
Internal Volume Response Time	The time it takes from the moment a request for information arrives at the storage device until the time when the storage device begins to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Internal Volume Response Time (R&W)	The time it takes from the moment a read or write request arrives at the storage device to the time when the storage device begins to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Top Internal Volume Response Time	The maximum time it takes from the moment a request for information arrives at the storage device until the time when the storage devices begin to send the information back in response.
Internal Volume IOPS	Measures the total number of I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Internal Volume IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top Internal Volume IOPS	The maximum number of I/O service requests generated by a single internal volume (across all storage internal-volumes) over the selected duration.
Internal Volume Throughput	Rate at which data was transferred to/from the internal volume in response to I/O service requests during the time period of the selected disk utilization violation (measured in MB per sec).

Internal Volume Throughput	Rate at which data was transferred to/from the internal volume in response to I/O service requests during the time period of the selected disk utilization violation (measured in MB per sec).
Internal Volume Throughput (R&W)	Amount of Read or Write data that is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec) for the internal volume.
Top Internal Volume Throughput	The maximum rate that data can be transmitted in a fixed amount of time to/from the internal volume.
Internal Volume Accessed	“Yes” indicates that the resource was accessed in the selected time range.
Disk Utilization	The percentage of post-cache service time used for requests out of the available sample time. This metric indicates what portion of time the disk is busy servicing requests.
Disk Utilization (R&W)	The number of Read or Write I/O service requests showing how the disk allocation is being used.
Top Disk Utilization	The maximum percentage of the total disk pre-cache Service Time generated by a single disk (across all storage disks) over the selected duration.
Disk IOPS	Measures the total number of I/O service requests on the disk for the virtual volume during the time period of the selected disk utilization violation (measured in I/O per sec).
Disk IOPS (R&W)	The number of Read or Write I/O service requests passing through the disk's I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top Disk IOPS	The maximum sum of IOPS reported by the measured disks.
Disk Throughput	Rate at which data was transferred to/from the disk in response to I/O service requests on the virtual volume during the time period of the selected disk utilization violation (measured in MB per sec).
Disk Throughput (R&W)	Amount of Read or Write data that is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec) for the disk.
Top Disk Throughput	The maximum rate that data can be transmitted in a fixed amount of time to/from the disk.
Disk Accessed	Indication if at least one disk was accessed (for read or write).
<i>icon (Is the device offline?)</i>	An icon in this column indicates that the device is offline.
<annotations>	User-defined terminology associated with the devices.

Related tasks

[Identifying performance trends](#) on page 30

[Analyzing virtual storage performance](#) on page 50

Storage Pool Performance view

With the Storage Pool Performance view, you can see how well storage virtualization is performing. A detailed comparison can be done by individually examining absolute front-end and backend values.

Navigation

From the OnCommand Insight Open menu, select **Performance > Storage Performance**. Click the **Storage Pool Performance** icon.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	The name of the storage pool.
Storage	The name of the storage system controlling the storage pool.
Application	The names of the applications using the storage pool.
Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
Hosts	Names of the hosts associated with the storage pool.
Type	The type of storage pool, for example, Aggregate for NetApp storage systems, RAID group, Thin Provisioning for a thin provisioned storage pool, or Backend Group for array virtualization.
Auto Tiering	A checkmark indicates that the selected storage pool is using the automatic storage tiering technology (for example, FAST VP).
<i>icon (Is virtual?)</i>	A "V" icon in this column indicates that the device is a virtualized volume.
Capacity (GB)	Size of the volume that is accessible to host applications, in gigabytes.
Used Capacity (GB)	The amount of capacity holding actual data in the storage pool. Includes usage based on all file types.

Used Capacity (%)	The percentage of capacity consumed in the storage pool in gigabytes.
Unused Capacity (GB)	The usable capacity that might be available for storing additional data on the storage pool in gigabytes.
Backend Capacity (GB)	The sum of the capacities of the backend storage used to store the storage pool. This column applies only to virtual storage pools, whose storage is allocated on storage arrays not directly controlled by the storage system identified in the Storage column.
Commit Ratio	The ratio of the total space on the storage pool to the capacity that is allocated from it. This value can be greater than 100% when thin provisioning is in use (the pool is overcommitted).
Response Time	The time it takes from the moment a request for information arrives at a storage device to the time when the storage device begins to send the information back in response, measured in milliseconds. The actual calculation of response time depends on the device vendor.
Response Time (R&W)	The time it takes for a read or write request to arrive at the storage device and to respond to the request.
Top Response Time	The maximum of the maximum response times reported by the measured devices.
IOPS	The rate at which input and output operations are directed at a storage device. Each operation might have a large or small amount of data transferred for it. The actual calculation of IOPS depends on the device vendor.
IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum sum of IOPS reported by the measured devices.
Throughput	The rate at which data is sent to a storage device in a fixed amount of time. The value is measured in megabytes per second. The actual calculation of throughput depends on the device vendor.
Throughput (R&W)	The rate at which data is read or written to a storage device in a fixed amount of time. The value is measured in megabytes per second. The actual calculation of throughput depends on the device vendor.
Top Throughput	The maximum sum of throughputs reported by the measured devices.
Backend Response Time	The time it takes from the moment a request for information arrives at the backend volumes to the time when the storage device begins to send the information back in response, measured in milliseconds. The actual calculation of response time depends on the device vendor. This column is only populated for virtual storage pools.

Backend Response Time (R&W)	The time it takes for a read or write request to arrive at the backend volumes and to respond to the request.
Top Backend Response Time	The maximum of the maximum response times reported by the backend volumes. This column is only populated for virtual storage pools.
Backend IOPS	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec). This column is only populated for virtual storage pools.
Backend IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top Backend IOPS	The maximum sum of IOPS reported by the backend volumes. This column is only populated for virtual storage pools.
Backend Throughput	The rate at which data is read or written to the backend volumes in a fixed amount of time. The value is measured in megabytes per second. The actual calculation of throughput depends on the device vendor. This column is only populated for virtual storage pools.
Backend Throughput (R&W)	The rate at which data is read or written to the backend volumes in a fixed amount of time. The value is measured in megabytes per second. The actual calculation of throughput depends on the device vendor.
Top Backend Throughput	The maximum sum of throughputs reported by the backend volumes. This column is only populated for virtual storage pools.
IOPS Efficiency	The ratio of IOPs to Backend IOPS. This indicates the efficiency gained by virtualization. For example, if there are 100 front-end (backend plus non-backend) IOPS and 25 backend IOPS, then the ratio is $(100 - 25)/100 = 75\%$. This means 75% of the I/O requests were handled within the virtualizer and 25% required I/O to backend storage. A negative number indicates that the overhead of virtualization (housekeeping operations) outweighs the gain of virtualization. This column is only populated for virtual storage pools.
IOPS Efficiency (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Throughput Efficiency	The ratio of Throughput to Backend Throughput. This column is only populated for virtual storage pools.
Throughput Efficiency (R&W)	The ratio of Throughput to Backend Throughput read and write requests. This column is only populated for virtual storage pools.
Accessed	"Yes" indicates that the resource was accessed in the selected time range.
Backend Accessed	"Yes" indicates that the backend resource was accessed in the selected time range.
<i>annotations</i>	The annotations associated with these devices.

Options

From the Storage Pools Performance detail view, right-click to show a pop-up menu containing the following options.

Analyze	Available only with the Perform license. Allows you to investigate the performance of the selected resources affected by the violation. For example, you can determine contention issues, availability issues, and array performance. The Storage Pool Summary tab provides information that might be needed for troubleshooting.
Analyze Contention	Available only with the Perform license. Allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are in contention for performance on the disk, a factor that can cause poor performance.
Analyze Storage Pools	Available only with the Assure license. Allows you to select a specific storage pool and assess its status related to the thin-provisioning policies. You can use this dialog box, instead of the Violations Browser, to see the current thin provisioning violations and how close the storage pool is to reaching the policy limits.
Show Disks	Shows all disks used by the selected storage pool in a separate view and provides the capacity information for the storage pools
Edit Annotations	Allows you to assign a note to this resource so that you can later group or filter the resources by the annotation. For example, you might want to group or filter resource by a specific note or tier.

Storage Resource Performance chart

This view enables you to see how well the storage resources that have already been allocated are performing.

Navigation

From the OnCommand Insight Open menu, select **Performance > Datastore Performance**. Select one or more lines in the view. Click the **Storage Resource Performance** icon.

Descriptions

Show Change Events	Select this option to display letters as change event markers on the chart. Clicking on the markers highlights the corresponding change events in the list. The change events on the list briefly describe what the change involved and when the change was detected by OnCommand Insight. Clicking on the description link opens the Changes detail view that gives more information about the changes that occurred at that point.
---------------------------	--

Pan	Click the hand icon and move the mouse pointer over the performance chart to display details for any point in the graph, and right-click to open the Analyze dialog box for a selected point.
Legend	Displays the metrics shown on the chart that include IOPS, throughput, and response time data. Click Change to select different metrics to display in the chart.
Time Range	<p>Changes the time period for the chart. You can select the following time ranges: Live Sample, Today, Last Hour, Last 24h, Last Week, All, and Custom.</p> <ul style="list-style-type: none"> • With the Custom range, you can select any time range up to the current time. • With the Live Sample, OnCommand Insight first retrieves data for the last hour, then asks the server for changes from the last update time to the current time as long the as the chart is displayed.
Date and Time Stamp on x axis	The x axis displays the date and time when OnCommand Insight sampled and retrieved the data.
Zoom	The y axis automatically scales to 5x greater than the greatest value in the chart. In addition, you can zoom into a portion of the chart using the Zoom tool. To zoom in, click the Zoom tool and click and drag across a section of the chart. A box defines the boundaries of the detail that you zoom into. To zoom out, hold down the Shift key and drag across a region. To reset the default zoom, click the Zoom to Fit tool. These settings persist across multiple viewings of the performance chart.
Export image or data	You can export the chart as a graphic file or export data from the chart.

Switch Port Performance view

The Switch Port Performance view displays traffic distribution data across the different devices on the storage network, such as switches, hosts, and storage devices at any given point in time. Every device is listed and can be sorted by data load as a percentage.

Use the Switch Port Performance view to:

- Evaluate traffic distribution.
- Evaluate port utilization.
- Identify resources that can be leveraged to improve utilization.
- Diagnose traffic bottlenecks.
- Diagnose congestion problems.
- Determine when to plan to add new capacity.
- Determine when to better utilize current capacity.

Navigation

From the OnCommand Insight Open menu, select **Performance > Switch Port Performance**.

Column descriptions

This view displays data in these columns.

<i>blank</i>	Applicable with any presentation order other than No Grouping. Column that organizes the data according to the selected grouping format (by, for example, device name, connected device name). The number in parentheses indicates the number of ports reported in each (grouped) row.
Switch	Name of the switch in the storage network.
Switch Port	Name of the switch port on the switch.
Has Errors	Number of errors for each host.
Value	The amount of traffic by the device indicated as being sorted by in the menu in the upper right corner.
Rx Value	Value received.
Tx Value	Value transmitted.
Distribution	Percentage of the total load carried by the path.
Fabric/VSAN	Address of fabric or virtual storage area.
VSAN	Name of Virtual Storage Area Network.
Speed	The physical speed (in gigabytes) of the switch port shown in this view.
Switch Blade	The physical, usually racked, chassis in which switch components reside. One switch entity may contain multiple blades. One blade may contain multiple controllers.
Controller	An I/O controller board located in the blade container. A controller may service multiple ports.
Transfer Rate	The average sampled transfer rate for the time period and metric selected in the view (measured in MBs).
Rx Transfer Rate	Receiving rate.
Tx Transfer Rate	Transmission rate.
Connected Device Type	Type of device in the storage network, for example, switch, host, or storage array.
Connected Device Name	Name of the device in the storage network.

Connected Device Port	Name of the device port on the storage network.
V-Host	Name of the host virtualization server.
Has CRC Errors	Number of CRC frame errors.
Sync Loss Count	Number of synchronization loss errors.
Signal Loss Count	Number of signal loss errors.
Class 3 Discards Count	The count of Fibre Channel Class 3 data transport discards.
Frame Too Short Count	The count of Fibre Channel data transmission frames that are too short.
Frame Too Long Count	The count of Fibre Channel data transmission frames that are too long.
BB Credit Errors	Fibre Channel uses buffer-to-buffer credits to control transmission flow. The credit value is decreased by one when a frame is sent and increased when a response is received. As the available credits for a given port approach zero, the error warns that the port will stop receiving transmissions when zero is reached and will not resume until the BB credits can be replenished.
Application	The application or service that resides on the host.
Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.

Options

The following options are available from the right-click menu:

Analyze	Available only with the Perform license. Allows you to investigate the performance of the selected resources affected by the violation. For example, you can determine contention issues, availability issues, and array performance. If automatic storage tiering might affect the data, only high-level information is displayed. For details, click the Auto tiering affects data link.
Analyze Contention	Available only with the Perform license. Allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are contending for performance on the disk, a factor that can cause poor performance.
Show Performance Histogram	Displays the relative amount of time the selected ports spent in the given utilization levels.

Show Daily Traffic Pattern	Identifies the flow of traffic by hour, day, week or by a custom calendar setting.
-----------------------------------	--

Related tasks

[Monitoring Fibre Channel switch performance](#) on page 23

[Selecting an active path for a virtual machine data store](#) on page 39

[Determining if a VM is affecting host performance](#) on page 41

Switch Threshold settings

You use the Switch Threshold settings to establish the minimum and maximum performance thresholds that constitute policies. When a threshold level is reached, an alert is issued.

Navigation

You can change the threshold settings with the **Tools > Settings > Switch Thresholds** option.

If you want to locate a potential problem with switch performance, select **Assurance > Switch Port Performance Alerts**. Right-click on a threshold shown in red and select the **Configure** option to make changes.

Setting options

To set a switch threshold, set the following parameters in the tabbed groups of thresholds:

- **Min:** The minimum level that can be reached without an alert being issued. If a level reaches less than this minimum for the specified time period, an alert is issued. To disable a threshold, clear the check box.
- **Max:** The maximum level that can be reached without an alert being issued. If a level exceeds this maximum for the specified time period, an alert is issued. To disable a threshold, clear the check box.
- **Period (minutes):** The amount of time during which the error must occur for an alert to be issued.

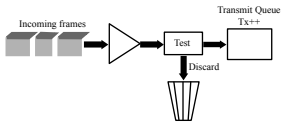
Performance threshold settings

Tx Utilization (%)	Percentage of available bandwidth used for sending transmissions.
Rx Utilization (%)	Percentage of available bandwidth used for receiving transmissions.
Utilization (%)	Percentage of available bandwidth used for Tx and Rx combined.
CRC (%)	The CRC frame errors as a percentage of the total data traffic. CRC frame errors indicate bit errors somewhere in the data path and point to poor connections, bad cables, or links that are too long.
Error (%)	The volume of errors as a percentage of the total data traffic.

Loss of Sync	If a Loss of Sync error occurs, the hardware cannot make sense of the traffic or lock onto it. All of the equipment might not be using the same data rate or the optics or physical connections might be of poor quality. The port must re-sync after each such error, which impacts system performance.
Loss of Signal	If a Loss of Signal error occurs, there is no electrical connection and a physical problem exists.
Class 3 Discards	The count of Fibre Channel Class 3 data transport discards.
Frame Size Too Long	The count of Fibre Channel data transmission frames that are too long.
Frame Size Too Short	The count of Fibre Channel data transmission frames that are too short.
BB Credit Errors	Fibre Channel uses buffer-to-buffer credits to control transmission flow. The credit value is decreased by one when a frame is sent and increased by one when a response is received. As the available credits for a given port approach zero, the error warns that the port will stop receiving transmissions when zero is reached and will not resume until the BB credits can be replenished.

Switch threshold types and formulas

OnCommand Insight uses these formulas to determine the switch port performance threshold data.

Threshold	Description	Formula
BB Credit Errors	Fibre Channel uses buffer-to-buffer credits to control transmission flow. The credit value is decremented when a frame is sent and replenished when a response is received. As the available credits for a given port approach zero, the error warns that the port will stop receiving transmissions when zero is reached and will not resume until the BB credits can be replenished.	
Class 3 Discards	The count of Fibre Channel Class 3 data transport discards.	
CRC Rate	CRC Rate is the measure of CRC frame errors as a percentage of the total data traffic. CRC frame errors indicate bit errors somewhere in the data path and point to poor connections, bad cables, or links that are too long.	$crcErrorRate = \frac{\Delta crcErrors}{\Delta RxFrames} \cdot 100$  <p>The diagram illustrates the process of CRC frame errors. It starts with 'Incoming frames' represented by three grey rectangles. These frames enter a trapezoidal 'Test' block. From the 'Test' block, one path leads to a 'Transmit Queue' (a rectangle labeled 'Tx++') and another path leads to a 'Discard' point, represented by a funnel icon.</p>

Threshold	Description	Formula
Errors Rate	Total number of errors (Loss of Sync, Loss of Signal, and Framing)	$errorRate(\%) = \frac{\Delta error \cdot 100}{\Delta error + \Delta RxFrames + \Delta TxFrames}$
Frame Size Too Long	The count of Fibre Channel data transmission frames that are too long.	
Frame Size Too Short	The count of Fibre Channel data transmission frames that are too short.	
Loss of Signal	If a Loss of Signal error occurs, there is no electrical connection and a physical problem exists.	$lossOfSignal / sec(\%) = \frac{\Delta lossOfSignal \cdot 100}{\Delta sec}$
Loss of Sync	If a Loss of Sync error occurs, the hardware cannot make sense of the traffic or lock onto it. All of the equipment might not be using the same data rate or the optics or physical connections might be of poor quality. The port must re-sync after each such error, which impacts system performance.	$lossOfSync / sec(\%) = \frac{\Delta lossOfSync \cdot 100}{\Delta sec}$
Received Utilization	Percentage of available bandwidth used for Rx.	$RxUtilization(\%) = \frac{\Delta RxBits \cdot 100}{\Delta Seconds \cdot ActualSpeedBits}$
Transmit Utilization	Percentage of available bandwidth used for Tx.	$TxUtilization(\%) = \frac{\Delta TxBits \cdot 100}{\Delta Seconds \cdot ActualSpeedBits}$
Utilization	Percentage of available bandwidth used for Tx and Rx.	$Utilization(\%) = \frac{\max(\Delta RxBits, \Delta TxBits) \cdot 100}{\Delta Seconds \cdot ActualSpeedBits}$

Topology view

Use this view to visualize your SAN or NAS environment, the devices, and their connections. Every device in your environment is shown as an icon that represents the device type, while physical

connections appear as lines connecting the devices. Each time you select a different device or path in a main view, the Topology representation changes as well.

Access

You can display a Topology map from the majority of the main views; however, you cannot access the Topology Map from the Switches, Storage Arrays, or Tapes main views. From a main view, select a device or path and click the Topology icon in the bottom of the Client view.

Operations

From the Topology view, you can perform the following operations:

- Position the mouse pointer over a device or path to see its detail.
- Click any device icon to view port information.
- Use the Topology toolbar to adjust settings.
- Select different view representations on the Topology Map.
- Change the link style to rounded or square.
- Add a watermark to the map.
- Add connected devices.
- Export the Topology layout as an image.

Related tasks

[Allocating capacity to an existing host](#) on page 62

Related references

[Analyze Contention dialog box](#) on page 68

[Analyze dialog box](#) on page 70

VMDK Performance view

Use this view to determine which disks are the top consumers of IOPS. From this list, you can choose the top consumers and use the VM Performance Distribution Chart to compare their IOPS.

Navigation

- From the OnCommand Insight Open menu, select **Performance > Virtual Machine Performance**. Click the VMDK Performance view icon.
- From the OnCommand Insight Open menu, select **Assurance > Disk Utilization Violations**. Select one or more virtual machines with disk utilization violations. Click the **Volume Usage of Disk** or **Internal Volume Usage of Disk** icon. Click the **VMDK Performance** icon.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format (by, for example, device name, connected device name). Applicable with any presentation order other than No Grouping. The number in parentheses indicates the number of ports reported in each (grouped) row.
Name	Name of the virtual machine disk.
Virtual Machine	Name of virtual machine.
Datastore	The name of data store residing on this virtual machine disk.
Capacity (GB)	Total storage array capacity that is accessible to host applications, in gigabytes.
Used Capacity (GB)	The amount of capacity holding actual data in the virtual machine disk. Includes usage based on all file types.
RDM	A VMware feature that exposes SCSI targets (or LUNs) directly to a virtual machine. RDMs are an alternative to using VMFS. RDMs are special files in a VMFS volume that act as a proxy for a raw device.
Host Names	Hosts associated with the virtual machine disk.
Storage	Storage arrays used by this virtual machine disk.
Resource Name	The volumes or internal volumes in the path.
Resource Technology	The SAN (FC and iSCSI) or NAS (NFS and CIFS) protocols that the device supports.
Resource Capacity (GB)	The total capacity, in gigabytes of the volumes or internal volumes, in the paths of this virtual machine's disk.
Resource Used Capacity (GB)	The total used capacity, in gigabytes of the volumes or internal volumes, in this virtual machine disk's paths.
Deduplication Savings	The known amount of storage savings through deduplication, a process that detects blocks with identical content and replaces subsequent identical blocks with a reference to a single copy of the block.
IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum sum of IOPS reported by the measured devices.

Throughput (R&W)	The rate at which data is read or written to the measured devices in a fixed amount of time. The value is measured in megabytes per second. The actual calculation of throughput depends on the device vendor.
Throughput	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Throughput	The maximum sum of throughputs reported by the measured devices.
Latency (R&W)	The rate at which data is read or written to the virtual machines in a fixed amount of time. The value is measured in megabytes per second.
Latency	The average response time from the virtual machines carved from a data store.
Top Latency	The highest response time from the virtual machines carved from a data store.

Options

The following options are available from the right-click menu:

Analyze	Available only with the Perform license. Allows you to investigate the performance of the selected virtual machines. For example, you can determine contention issues, availability issues, and disk performance.
Analyze Contention	Available only with the Perform license. Allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are in contention for performance on the disk, a factor that can cause poor performance.

VM Distribution view

This view lists the virtual machines using the selected virtual machine (VM) in their data path and their estimated I/O traffic distribution on the VM.

Navigation

From the OnCommand Insight Open menu, select the **Switch Port Performance** view, select a virtual machine, and click the **VM Distribution** icon.

Column descriptions

VM	The name of the selected virtual machine.
Datastore	The name of data store used by this virtual machine.
Host Port	The name of the port used on the host by this virtual machine.

Storage Port	The name of the port the virtual machine uses on the storage device.
Volume	The name of the volume this virtual machine uses in the storage device.
I/O	The total throughput of this virtual machine.
I/O Distribution	The estimated percentage contribution of this particular virtual machine to the total traffic on this port originating from any virtual machine.

Virtual Internal Volume Performance view

You use this view to see the details of the virtual internal volume involved in a disk utilization violation.

Navigation

1. From the OnCommand Insight Open menu, select **Assurance > Disk Utilization Violations**.
2. Select one or more items in the view. Click the **Volume Usage of Disk** icon to display a center view.
3. Click the **Virtual Internal Volume Performance** icon.

Alternately, right-click on any Performance view and select **Analyze Contention**. The Disks tab has Volume Usage and Internal Volume Usage detail views from which you select the **Virtual Internal Volume Performance** icon.

Column descriptions

<i>blank</i>	Applicable with any presentation order other than No Grouping. Column that organizes the data according to the selected grouping format.
Name	Name of the virtual volume.
Storage	Name of storage array on which the internal volume resides.
vFiler	The name of the vFiler unit. A vFiler unit is an isolated software container that behaves exactly like a physical storage array. A vFiler unit shares the physical resources of the array, but abstracts the client access from the physical array into virtual arrays.
Application	Applications associated with this virtual volume.
Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.

Hosts	Hosts making use of the virtual internal volume.
Capacity (GB)	Total storage array capacity that is accessible to host applications, in gigabytes.
Raw Capacity (GB)	Physical disk capacity of the volume, in gigabytes. This differs from usable capacity when technologies such as RAID-5 are used, where some of the raw capacity is used for protection purposes.
Used Capacity (GB)	The amount of capacity holding actual data in the virtual machine disk. Includes usage based on all file types.
Consumed Capacity	The amount of capacity that the volume consumes from underlying storage. For non-thin provisioned value, that is the same as the volume's capacity. For thin provisioned volumes, it is the amount of capacity used to store the volume's contents. Its value is also affected by Snapshot copies, deduplication, and other storage technologies. If the volume does not have usage information available, this appears blank.
Storage Pool	Storage arrays used by this virtual machine disk.
Datastore	The name of data store residing on this internal volume.
Type	Type of internal volume, for example, FlexClone or FlexVol.
Thin Provisioned	A check mark indicates that the resource is thin provisioned.
Performance Policy	The level of performance threshold (for example, Global) set on this resource. For example, maximum IOPS, response time, or throughput thresholds might be set for specific resources and not use the global thresholds.
Space Guarantee	A vendor-specific indication of how the internal volume's space is allocated if thin provisioning is used.
Deduplication Savings	Percentage value of the rate of deduplication in effect for the internal volume.
Clone Source	The name of the internal volume that this internal volume cloned. This is the source of the cloned relationship.
Clone Shared Capacity (GB)	When the Internal Volume is a clone of another Internal Volume, this value shows the amount of capacity that is shared between the Source and Target clones (If no changes have been made to the source and target Internal Volumes since the clone operation occurred, all capacity will be shared. Modifications made to either of the Internal Volumes will decrease the shared capacity).
Status	Information about whether the internal volume is online, offline, or other status.

Snapshot Reserve (GB)	The capacity that was reserved for Snapshot copies.
Snapshot Used (GB)	The capacity remaining after some capacity that was reserved for Snapshot copies was actually used.
Snapshot Used (%)	The percent of capacity remaining after some capacity that was reserved for Snapshot copies was actually used.
Snapshot Overflow (GB)	The capacity that was used for Snapshots, but exceeded what was reserved for Snapshot copies.
Snapshot Count	The number of Snapshot copies that are stored for this internal volume.
Last Snapshot	The time when the last Snapshot copy operation occurred on this internal volume.
Disk Types	The types of physical disks (for example, Fibre Channel or ATA) on which the internal volume is based.
Disk Size (GB)	The size of the physical disks on which the internal volume is based.
Disk Speed (RPM)	The speed (in revolutions per minute) of the physical disks on which the internal volume is based.
Response Time	The time in milliseconds that it took to access the disk. These values are based on fixed speed specifications of the disks.
Response Time (R&W)	The time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.
Top Response Time	The maximum length of time in milliseconds that it took for the internal volume to begin to send the response.
IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).
IOPS (R&W)	The number of Read and Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum number of I/O service requests generated by the virtual internal volume over the selected length of time.
Throughput	Rate at which Read or Write data can be received in a fixed amount of time in response to I/O service requests (measured in MB per sec).

Throughput (R&W)	Rate at which data is read or written in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Throughput	Maximum rate at which Read or Write data can be received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Accessed	“Yes” indicates that the resource was accessed in the selected time range.
<i>annotations</i>	User-defined terminology associated with each virtual internal volume including Service Level, Note, and Tier.

Virtual Machine Performance view

Use this view to look for inactive VMs or VMs that are using high storage that could be using lower-tier storage.

Navigation

From the OnCommand Insight Open menu, select **Performance > Virtual Machine Performance**.

Column descriptions

It might be helpful to group the data by data stores. One VM can access multiple data stores.

With this view, you can select a VM and see additional data in other views:

- Select multiple VMs and with the VM Performance Distribution detail view, you can compare traffic distribution and identify which VM is generating the most IOPS.
- With the Virtual Machines Performance Chart, you can see how the load (CPU or memory) is distributed across the VMs on a specific ESX host and see if there are sudden changes or anomalies.

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Name	Name of virtual machine. This view does not display information about virtual machines and LUNs that do not map to a volume.
DNS Name	The host domain name.
Host Names	Names of hosts using the virtual machine.
Datastore	The name of data store using the virtual machine.
V-Cluster	Name of the cluster.
CPU Utilization	Amount of actively used CPU, as a percentage of total available (over all virtual CPUs).

Memory Utilization	Threshold for the memory used by the host.
Disk Latency	The sum of Disk Read Latency and Disk Write Latency for the sampling period.
Top Disk Latency	The maximum (over all virtual disks) amount of time taken for a virtual disk write operation (ms).
Disk Read Latency	Average amount of time taken during the interval to process a SCSI read command issued from the Guest OS to the virtual machine. Latency data is derived from the VMDK.
Disk Write Latency	Average amount of time taken during the interval to process a SCSI write command issued by the Guest OS to the virtual machine. The average (over all virtual disks) amount of time taken for a virtual disk write operation (ms).
Disk IOPS	The sum of Disk Read IOPS and Disk Write IOPS.
Top Disk IOPS	The maximum sum of IOPS reported by the measured disks.
Disk Read IOPS	The number of read I/O service requests passing through the disk's I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Disk Write IOPS	Number of disk write operations (across all disks) per second in the latest poll.
Disk Throughput	Aggregated disk I/O rate (MB/s).
Top Disk Throughput	The maximum rate that data can be transmitted in a fixed amount of time to/from the disk.
Disk Read Throughput	Total averaged rate to read from the disk in megabytes per second.
Disk Write Throughput	Total averaged rate written to the disk in megabytes per second.
IP Throughput	Aggregated rate at which IP data was transmitted and received in megabytes per second..
IP Throughput (Receive)	Average rate at which IP data was received in megabytes per second.
IP Throughput (Transmit)	Average rate at which IP data was transmitted in megabytes per second.
Disk Accessed	Indication if at least one virtual disk was accessed (for read or write).
IP Accessed	Indication if Guest OS received or transmitted any IP data.

Options

The following options are available from the right-click menu:

Analyze	Available only with the Perform license. Allows you to investigate the performance of the selected resources affected by the violation. For example, you can determine contention issues, availability issues, and array performance. The Virtual Machine Summary tab provides information that might be needed for troubleshooting.
Analyze Contention	Available only with the Perform license. Allows you to analyze the reasons for poor performance and identify where that performance has degraded. For example, you can see whether hosts and applications are in contention for performance on the disk, a factor that can cause poor performance.
Analyze Storage Pools	Available only with the Assure license. Allows you to select a specific storage pool and assess its status related to the thin provisioning policies. You can use this dialog box, instead of the Violations Browser, to see the current thin provisioning violations and how close the storage pool is to reaching the policy limits.

Related tasks

[Determining virtual machine contention](#) on page 37

[Viewing virtual machine performance and utilization](#) on page 34

Related references

[Virtual Machine Summary tab](#) on page 78

Virtual Volume Usage view

You use this view to see the details of the virtual machine involved in a disk utilization violation.

Navigation

From the OnCommand Insight Open menu, select **Assurance > Disk Utilization Violations**. Select one or more items in the view. Click the **Volume Usage of Disk** icon. Click the **Virtual Volume Usage** icon.

Column descriptions

blank	Applicable with any presentation order other than No Grouping. Column that organizes the data according to the selected grouping format.
Volume	Name of the virtual volume.
Label	Full virtual volume name within the storage array.

Type	The vendor-specific type of virtual volume (for example, SFS or B.V.).
<i>array virtualization type icon</i>	Indicates type of virtualization. Showing on a virtual volume, a "V" icon indicates that the device is a virtualized volume and a "B" icon indicates that the device is a backend volume.
Hosts	Hosts making use of the virtual volume.
Application	Applications associated with this virtual volume.
Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
Volume IOPS	Measures the total number of I/O service requests on the virtual volume during the time period of the selected disk utilization violation (measured in I/O per sec).
Disk IOPS	Measures the total number of I/O service requests on the disk for the virtual volume during the time period of the selected disk utilization violation (measured in I/O per sec).
Disk IOPS Distribution	Portion of the disk IOPS that the virtual volume contributed toward the total Disk IOPS during the time period of the selected disk utilization violation.
Volume Throughput	Rate at which data was transferred to or from the virtual volume in response to I/O service requests during the time period of the selected disk utilization violation (measured in MB per sec).
Disk Throughput	Rate at which data was transferred to/from the disk in response to I/O service requests on the virtual volume during the time period of the selected disk utilization violation (measured in MB per sec).
Disk Throughput Distribution	Portion of the disk throughput that the virtual volume contributed toward the total disk throughput during the time period of the selected disk utilization violation.
Volume Response Time	The time in milliseconds that it takes from the moment a request for information arrives at the storage device to the time when that storage device begins to send information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array.
Volume Top Response Time	The maximum time in milliseconds it takes to send the information back.

Volume Performance view

Using the Volume Performance view, you can see all the volumes that the hosts are using and the metrics necessary to analyze volume performance. This data helps you better use capacity and assists in proactive analysis of how the storage resources are being used. For SAN environments, the Volume Performance view shows the end points of the SAN paths.

Navigation

- From the OnCommand Insight Open menu, select **Performance** and the Host, Storage, Application, Datastore, or Virtual Machine performance views. At the bottom of the any of these views, click the **Volume Performance** icon.
- From one of the Performance views, select a row, right-click, and select Analyze. Click the **Volumes** tab.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
Volume	Name of volume.
Label	An alternate name or alias assigned to a volume by the storage administrator.
UUID	Universally unique identifier for the object. In this case, it is generated by and retrieved from the storage array itself.
Storage	Name of the storage array.
Storage Pool	The name of the storage pool on which the volume resides.
Auto Tiering	A checkmark indicates that the selected storage pool is using the automatic storage tiering technology (for example, FAST VP).
Thin Provisioned	A checkmark indicates that the volume is leveraging thin provisioning.
Performance Policy	The level of performance threshold (for example, Global) set on this resource. For example, maximum IOPS, response time, or throughput thresholds might be set for specific resources and not use the global thresholds.
vFiler	The name of the vFiler unit. A vFiler unit is an isolated software container that behaves exactly like a physical storage array. A vFiler unit shares the physical resources of the array, but abstracts the client access from the physical array into virtual arrays.
Internal Volume	Name of the internal volume that the volume uses.
Qtree	Name of the qtree on this volume.

<i>array virtualization type icon</i>	Indicates type of virtualization. Showing on a virtual volume, a "V" icon indicates that the device is a virtualized volume and a "B" icon indicates that the device is a backend volume.
Virtualizer	For backend volumes in array virtualization. Displays the name of the front end virtualizer that is using this volume.
Virtual Storage Pool	For backend volumes in array virtualization. The name of the storage pool on the front end virtualizer that is using this volume.
Datastore	The name of data store residing on this volume.
Application	Application associated with this volume.
Application Priority	The importance of this application within your organization that was set when defining the application.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
Hosts	Hosts making use of the volume.
Capacity (GB)	Size of the volume that is accessible to host applications, in gigabytes.
Raw Capacity (GB)	Physical disk capacity of the volume, in gigabytes. This differs from usable capacity when technologies such as RAID-5 are used, where some of the raw capacity is used for protection purposes.
Consumed Capacity (GB)	The amount of capacity that the volume consumes from underlying storage (for example, internal volume and storage pool). For non-thin provisioned volumes, this value is the same as the volume's capacity. For thin provisioned volumes, it is the amount of capacity used to store the volume's contents. Its value is also affected by Snapshot copies, deduplication, and other storage technologies. If the volume does not have usage information available, this appears blank.
Mapped Ports	The number of storage ports through which this volume is accessible.
Redundancy	Level of mirroring defined for the device based on the storage technology, for example, RAID-DP, underlying the device. This is taken from the device itself. For an explanation of values, see the device documentation.
Response Time	The time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response.
Response Time (R&W)	The time it takes for a read or write request to arrive at the storage device and to respond to the request.

Top Response Time	The maximum length of time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response.
IOPS	The portion or ratio of I/O service requests by the selected host or application passing through the I/O channel per unit of time (measured in I/O per sec).
IOPS (R&W)	The number of Read or Write I/O service requests passing through the I/O channel or portion of that channel per unit of time (measured in I/O per sec).
Top IOPS	The maximum number of I/O service requests generated by the volume over the selected length of time.
Throughput	Rate that data is being transmitted in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Throughput (R&W)	Rate at which Read or Write data is being received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Top Throughput	Maximum rate at which data can be received in a fixed amount of time in response to I/O service requests (measured in MB per sec).
Cache Hit Ratio	Percentage of requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better the performance. This column is empty for storage arrays that do not collect cache hit information.
Cache Hit Ratio (R&W)	Percentage of Read/Write requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better the performance. This column is empty for storage arrays that do not collect cache hit information.
Partial R/W	Total number of times that a read or write crosses a stripe boundary on any disk module in RAID 5, RAID 1/0 or RAID 0 LUN. Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size. A high percentage is often an indication of improper alignment of a volume (or a NetApp LUN). For CLARiiON, this is the number of stripe crossings divided by the total number of IOPS.
Write Pending	The number of Write I/O service requests that are pending.
Accessed	"Yes" indicates that the resource was accessed in the selected time range.
<i>annotations</i>	Annotations associated with each volume.

Options

From the Volume Performance detail view, right-click to display a menu containing the following options:

Analyze	Allows you to identify potential problems by isolating the location of contention and the source of delays and by comparing performance data for different time periods to note changes. The Volume Summary tab provides information that might be needed for troubleshooting.
Analyze Contention	Allows you to analyze the reasons for poor performance and identify where that performance has degraded.
Analyze Storage Pools	Allows you to select a specific storage pool and assess its status related to the thin-provisioning policies. You can use this dialog box, instead of the Violations Browser, to see the current thin provisioning violations and how close the storage pool is to reaching the policy limits.
Show Disks	Shows all disks used by the selected storage pool in a separate view.
Show Volume Members	Lists the members of the volume in a separate window.
Modify Policy	Changes the policy that governs alerts set on this resource. This option requires the Assure license.
Edit/Clear/Set Annotations	Allows you to assign a note to this resource so that you can later group or filter the resources by the annotation. For example, you might want to group or filter resource by a specific note or tier. You can also change annotations or remove them.

Related tasks

[Allocating capacity to a new host](#) on page 60

Related references

[Volume Summary tab](#) on page 79

Volume Usage of Disk view

To examine additional information for each volume in the Disk Utilization Violations view, open the Volume Usage of Disk view.

Navigation

From the OnCommand Insight Open menu, select **Assurance > Disk Utilization Violations**. At the bottom of the view, click the **Volume Usage of Disk** icon.

Column descriptions

<i>blank</i>	Column that organizes the data according to the selected grouping format. Applicable with any presentation order other than No Grouping.
--------------	--

Volume	Name of the volume.
Label	Full volume name within the storage array.
Type	The vendor-specific type of volume (for example, SFS or B.V.).
<i>icon (array virtualization type)</i>	Indicates type of virtualization. A "V" icon showing on a volume indicates that the device is a virtualized volume, and a "B" icon showing on a volume indicates that the device is a backend volume.
Hosts	Hosts using the volume.
Application	Applications associated with this volume.
Application Priority	Importance of this application as defined when establishing the business entities.
Tenant, Line of Business, Business Unit, Project	Columns listing the business entity components associated with the applications.
Volume IOPS	Measures the total number of I/O service requests on the volume during the time period of the selected disk utilization violation (measured in I/O per sec).
Disk IOPS	Measures the total number of I/O service requests on the disk for the volume during the time period of the selected disk utilization violation (measured in I/O per sec).
Disk IOPS Distribution	Portion of the disk IOPS that the volume contributed toward the total Disk IOPS during the time period of the selected disk utilization violation.
Volume Throughput	Rate at which data was transferred to/from the volume in response to I/O service requests during the time period of the selected disk utilization violation (measured in MB per sec).
Disk Throughput	Rate at which data was transferred to/from the disk in response to I/O service requests on the volume during the time period of the selected disk utilization violation (measured in MB per sec).
Disk Throughput Distribution	Portion of the disk throughput that the volume contributed toward the total disk throughput during the time period of the selected disk utilization violation.
Volume Response Time	The time it takes from the moment a request for information arrives at the storage device to the time when the storage devices begin to send the information back in response. This is the actual latency of the device, based on live samplings whenever possible. OnCommand Insight measures response times from within the storage array. The value is measured in milliseconds.

Volume Top Response Time The maximum time it takes to send the information back.

Options

From the Volume Usage of Disk detail view, right-click to show a menu containing the following options:

- | | |
|------------------------------|--|
| Analyze | Allows you to identify potential problems by isolating the location of contention and the source of delays and by comparing performance data for different time periods to note changes. The Volume Summary tab provides information that might be needed for troubleshooting. |
| Analyze Contention | Allows you to analyze the reasons for poor performance and identify where that performance has degraded. |
| Analyze Storage Pools | Allows you to select a specific storage pool and assess its status related to the thin provisioning policies. You can use this dialog box, instead of the Violations Browser, to see the current thin provisioning violations and how close the storage pool is to reaching the policy limits. |

Copyright information

Copyright © 1994–2013 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and www.ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- active path for virtual machines [39](#)
- alerts
 - adjusting switch thresholds [11](#)
 - disabling performance thresholds [14](#)
 - performance [14](#)
 - reviewing and confirming [15](#)
- Analyze Contention dialog box [68](#)
- analyzing
 - performance [41](#), [43](#), [126](#)
 - port balance [53](#)
 - traffic flow [46](#)
- applications
 - assigning for efficiency [58](#)
 - performance [44](#), [81](#)
- Assure
 - features [7](#)
- automatic storage tiering
 - in performance views [67](#)

B

- backend volumes
 - performance [83](#)
- balancing
 - data flow across ISLs [59](#)
 - data flow across ports [59](#)
 - ports [53](#), [55](#)
 - resources [53](#), [56](#)

C

- cabling
 - monitoring changes to [21](#)
- capacity
 - allocating to existing host [62](#)
 - allocating to new host [60](#)
 - identifying volumes with available [57](#)
- change markers
 - displaying [101](#)
 - linked to Changes view [20](#)
- changes
 - markers [50](#)
 - monitoring [21](#)
- configuration

- monitoring changes to [21](#)
- congestion
 - identifying [46](#), [48](#)
 - identifying cause [48](#)
 - storage port [49](#)
- contention
 - analyzing [68](#)
 - virtual machines [37](#)
- CPU utilization
 - threshold [35](#)
 - virtual machines [37](#)

D

- Daily Performance Pattern
 - histogram [71](#)
 - trend analysis [18](#)
- data
 - exporting performance charts [32](#)
 - grouping by Data Center annotation [60](#)
 - grouping in views [32](#)
 - migration to different tiers [58](#)
- data stores
 - adding for load balancing [36](#)
 - performance [72](#), [86](#)
 - performance analysis [24](#)
 - sharing in cluster [40](#)
- devices
 - monitoring changes [21](#)
 - performance [16](#)
 - switch performance [23](#)
- disk usage
 - internal volumes [100](#)
- disks
 - performance [89](#)
- Distribution Chart
 - export as image or file [32](#)
 - trend analysis [18](#)
- documentation
 - list of [9](#)
 - videos [10](#)

E

- examples
 - host performance troubleshooting [41](#)

- load balancing [36](#)

- searching [16](#), [41](#)

- exporting

- performance data [32](#)

F

- Fibre Channel

- switch performance [23](#)

G

- global policies

- customizing general [27](#)

- hierarchies for customization [29](#)

- grouping data

- by annotations [60](#)

H

- hosts

- adding to virtual cluster [40](#)

- allocating capacity to existing [62](#)

- allocating capacity to new [60](#)

- analyzing performance [43](#)

- identifying congested [48](#)

- performance [75](#), [93](#)

- performance analysis [16](#)

- troubleshooting performance [41](#)

- virtual machine utilization thresholds [105](#)

- virtualization [105](#)

I

- I/O utilization

- virtual machines [37](#)

- internal volumes

- disk usage [100](#)

- performance [76](#), [97](#)

- ISLs

- balancing data flow across [59](#)

L

- latency

- threshold [35](#)

M

- memory

- utilization threshold [35](#)

- monitoring

- storage performance [25](#)

- system changes [21](#)

O

- OnCommand Insight

- product portfolio [7](#)

- outages

- monitoring changes to [21](#)

P

- paths

- setting for virtual machine data store [39](#)

- Perform

- features [7](#)

- license required [70](#)

- performance

- alerts [14](#)

- analyzing [43](#), [70](#)

- applications [44](#), [81](#)

- automatic storage tiering effects [67](#)

- backend volumes [83](#)

- change markers [20](#)

- chart [18](#), [20](#), [30](#), [101](#)

- customizing general policies [27](#)

- daily pattern histogram [71](#)

- data stores [24](#), [72](#), [86](#)

- disabling thresholds [14](#)

- disks [89](#)

- distribution [103](#)

- examining [16](#)

- exporting data [32](#)

- hosts [75](#), [93](#)

- internal volumes [76](#), [97](#)

- monitoring switches [23](#)

- report [65](#)

- reports [64](#)

- searching data [16](#), [41](#)

- storage [25](#)

- storage array [106](#)

- storage pools [77](#)

- switch ports [114](#)

- system changes impacting [21](#)

- trend analysis [18](#)

- trends [30](#)
- troubleshooting [43](#)
- troubleshooting on hosts [41](#)
- utilization bar chart [92](#)
- virtual internal volumes [123](#)
- virtual machines [34](#), [78](#), [126](#)
- virtual storage [50](#)
- VM disks [120](#)
- volumes [79](#), [130](#)

Performance Chart

- change markers [20](#)
- export as image or file [32](#)
- trend analysis [18](#)

Plan

- features [7](#)

policies

- customizing global general [27](#)
- global hierarchy [29](#)
- global performance [16](#)
- host virtualization [105](#)
- monitoring changes to [21](#)
- virtual machine host thresholds [105](#)

ports

- balance violations [53](#), [55](#)
- balancing data flow [59](#)
- congested storage [49](#)
- in a cluster [75](#)
- mapped to volumes [79](#)

product documentation

- list of [9](#)

R

reports

- Array Performance [65](#)
- performance [64](#)

S

searching

- performance data [16](#), [41](#)

Settings

- virtual machine thresholds [35](#)

storage

- allocation [56](#)
- performance [25](#), [106](#)
- port balancing [56](#)
- port congestion [49](#)

storage pools

- performance [77](#)

switch ports

- performance [114](#)
- traffic flow [46](#)

switch thresholds

- adjusting for policies [11](#)
- types and formulas [12](#), [118](#)

switches

- changing thresholds [14](#)
- performance [23](#)
- performance alerts [14](#)

system changes

- monitoring [21](#)

T

thresholds

- array performance [25](#)
- global performance [16](#)
- over performance policy [43](#)
- performance [14](#)
- setting for virtual machine host utilization [105](#)
- virtual machines [35](#)

tiers

- migration research [58](#)

topology

- locating all connected devices [62](#)
- view operations [119](#)

traffic

- analyzing flow [46](#)
- daily pattern [71](#)
- distribution illustrated [103](#)

trend analysis

- histogram bar chart [92](#)
- performance [30](#)
- Performance Chart [101](#)
- Performance Distribution [103](#)
- tools [18](#)

troubleshooting

- poor application performance [44](#)
- slow client computers [43](#)
- VM performance [41](#)

U

utilization

- setting thresholds on hosts [105](#)
- virtual machine [34](#)

V

videos

- Community portal [10](#)
- performance visibility [16](#)
- troubleshooting performance [43](#)

views

- grouping data [32](#)

violations

- general [27](#)
- port balance [53](#)
- storage port balance [56](#)

virtual machine hosts

- setting utilization thresholds [105](#)

virtual machines

- adding [36](#)
- contention [37](#)
- disk performance [120](#)
- performance [34](#), [78](#), [126](#)
- searching [41](#)

- setting active path for [39](#)

- setting thresholds [35](#)

- traffic distribution [122](#)

virtual storage

- analyzing performance [50](#)

volumes

- backend [83](#)
- disks used list [133](#)
- internal [97](#)
- orphaned [65](#)
- performance [79](#), [130](#)
- performance analysis [16](#)
- tiered [65](#)
- with potential capacity [57](#)

Z

zoning

- monitoring changes to [21](#)