

Hitachi Ops Center Analyzer

10.8.2

User Guide

© 2016, 2022 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at <https://support.hitachivantara.com/en-us/contact-us.html>.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Edge, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.

Contents

Preface	9
Product version.....	9
Intended audience.....	9
Release notes.....	9
Related documents.....	9
Document conventions.....	9
Conventions for storage capacity values.....	11
Accessing product documentation.....	12
Getting help.....	12
Comments.....	12
Chapter 1: Ops Center Analyzer overview	13
Product overview.....	13
Features of Ops Center Analyzer.....	14
Unified infrastructure monitoring dashboard.....	14
Advanced reporting.....	15
SLO management.....	16
End-to-end monitoring.....	16
Problem identification and root cause analysis.....	17
Storage I/O controls.....	17
About Hitachi Virtual Storage Software Block (VSSB).....	18
Hitachi Ops Center Automator integration.....	18
Risk management capability.....	18
Getting started in your Ops Center Analyzer environment.....	18
Logging on to Ops Center Analyzer	18
Accessing Ops Center Analyzer detail view.....	19
Launching Ops Center Automator from Ops Center Analyzer.....	20
Setting up a monitoring environment for your infrastructure resources.....	20
Chapter 2: Quick access to vital performance reports	22
Dashboard overview.....	22
System Status Summary for Consumers and User Resources.....	24
System Status Summary for Consumers.....	24
System Status Summary for User Resources.....	25
System Resource Status report.....	25

Event Trends report.....	26
Resource Events report.....	27
System and Resource Events.....	27
Sending dashboard reports to users.....	28
Customizing the dashboard.....	29
Chapter 3: Creating custom reports.....	31
Custom reports overview.....	31
Creating custom reports.....	32
Managing custom reports.....	34
Examples of creating custom reports.....	34
Creating a line graph.....	35
Creating a table report.....	38
Sample queries.....	40
Chapter 4: Analyzing performance problems.....	41
Identifying performance problems.....	41
Infrastructure components and key performance metrics.....	42
Analytics workflow.....	44
Detecting performance problems.....	44
Analyzing performance bottlenecks.....	45
Identifying the bottleneck in E2E view.....	46
Comparing performance trends in Sparkline view.....	49
Identifying affected resources.....	50
Analyzing the cause of the bottleneck.....	50
Analyzing shared resources.....	51
Analyzing configuration changes.....	52
Checking recovery plans.....	53
Executing actions.....	54
Submitting services through Ops Center Automator.....	56
Granular data collection.....	56
Collecting granular data.....	57
Using Hitachi Ops Center Analyzer for data analysis: from deep dive to recovery planning.....	60
Chapter 5: Analyzing performance trends with predictive analytics risk reporting.....	69
Predictive analytics risk reporting overview.....	69
Anticipate performance trends	69
About predictive risk profiles.....	70
About predictive risk report definitions.....	70
Predictive analytics risk reporting workflow.....	71
Adding the Predictive Analytics license.....	71

Creating a predictive risk profile.....	72
Creating risk report definitions.....	72
Generating a risk report.....	73
Understanding performance trend projections in risk reports.....	74
Enhancing infrastructure management capabilities.....	75
Periodic analysis with predictive analytics risk reporting.....	78
E2E View analysis with predictive analytics risk reporting.....	80
Preventive actions and predictive analytics risk reporting.....	82
Chapter 6: Optimizing infrastructure resources with storage I/O controls.....	84
I/O control overview.....	84
I/O control operations.....	84
Requirements.....	85
Search capabilities for target volumes.....	85
Upper limit setting metrics.....	85
Use cases for setting upper limits.....	85
I/O control settings for an SLO.....	86
I/O controls for optimizing infrastructure resources.....	86
I/O controls for optimizing performance after the bottleneck analysis.....	86
Performing storage I/O control tasks in Ops Center Analyzer with Ops Center Automator.....	88
Setting I/O control limits.....	88
Clearing I/O control limits.....	89
Chapter 7: Configuring resource monitoring.....	90
Overview of managing and monitoring infrastructure resources.....	90
Resource monitoring settings.....	91
Dynamic thresholds.....	92
Advantages of dynamic thresholds.....	92
Determining if the computed value is correct.....	92
Automatic calculation of baseline values.....	92
Monitoring using a dynamic threshold.....	94
Dynamic threshold monitoring margins.....	95
Selecting the dynamic threshold margin.....	95
Static thresholds.....	96
Setting static thresholds using monitoring profiles.....	97
.....	98
Monitoring using a static threshold.....	99
About default profiles for volumes (user) and arrays (system).....	99
Setting thresholds for user resources.....	103
Creating a user resource assignment rule.....	104

Changing user resource assignment rule priority.....	105
Running a user resource assignment rule.....	105
Setting thresholds for system resources.....	105
Removing monitored resources.....	106
Chapter 8: Managing consumers.....	109
Consumer settings.....	109
Creating a consumer.....	110
Creating a resource assignment rule.....	110
Changing a resource assignment rule priority.....	111
Running a resource assignment rule.....	111
Creating multiple consumers (batch mode).....	112
About the consumer definition file format.....	113
Chapter 9: Setting notifications.....	114
Email notification settings.....	114
Configuring the mail server.....	114
Creating a condition profile.....	115
Enabling or disabling email addresses.....	115
Chapter 10: Managing users.....	116
User management.....	116
User permissions for Hitachi Ops Center Automator services	118
Configuring external authentication for users.....	118
Configuring external authentication for groups.....	119
Security.....	120
Chapter 11: Additional dashboard reports.....	121
System Status Summary for Storage Resources.....	121
System Status Summary for Volumes.....	121
System Status Summary for System Resources.....	122
System Status Summary for Infrastructure Resources.....	122
System Status Summary for User Resources.....	122
System Resource Status report.....	123
Consumer reports.....	124
Consumer Summary report.....	124
Consumers report.....	124
Consumers - Critical report.....	125
Consumers - Critical Each Grade report.....	125
VM reports.....	126
VM Summary report.....	126
VMs report.....	127
VMs - Critical.....	127

VM CPU Ready report.....	128
VM NIC Dropped report.....	128
VM Disk Latency report.....	129
Total VMs report.....	129
Volume reports.....	130
Volume Summary report.....	130
Volumes report.....	130
Volumes - Critical report.....	131
Volume IO Rate report.....	131
Volume IOPS report.....	132
Volume Response Time report.....	132
Total Volumes report.....	133
Host reports.....	133
Hosts report.....	133
Hosts - Critical report.....	133
Total Hosts report.....	134
User resource reports.....	134
Total User Resources report.....	134
Hypervisor reports.....	135
Hypervisors report.....	135
Total Hypervisors report.....	136
Switch reports.....	136
Switches report.....	136
Total Switches report.....	137
Storage system reports.....	137
Storage Systems report.....	137
Total Storage Systems report.....	138
System resource reports.....	138
System Resource Summary report.....	138
System Resources - Critical report.....	139
System Resources report.....	139
Total System Resources report.....	139
Event reports.....	140
System Event Status report.....	140
Resource Event Status report.....	140
Capacity reports.....	141
Capacity of VMware Datastores report.....	141
Capacity of VMware Datastores with Usage Exceeding 80% report.....	141
Capacity by Storage System report.....	142
Capacity by Storage System report (HUS 100 Storage Systems).....	143
Saving Ratio by Storage System report.....	143

Total Efficiency by Storage System report.....	144
Capacity of Drive Types by Storage System report (HDP).....	145
Capacity of Drive Types by Storage System report (HDT).....	147
Capacity of Drive Types by Storage System report (HUS 100 Storage Systems).....	149
Capacity by Pool report.....	151
Capacity by Pool report (HUS 100 Storage Systems).....	152
Total Efficiency by Pool report.....	153
Capacity by Consumer report.....	154
Capacity by Consumer in the Past 6 Months report.....	154
Capacity by Consumer report (HUS 100 Storage Systems).....	155
Capacity by Consumer in the Past 6 Months report (HUS 100 Storage Systems).....	155
Chapter 12: Performance analytics and best practices.....	157
Understanding utilization.....	157
Utilization metrics.....	158
Performance analysis.....	159
Performance analysis example.....	160
Performance troubleshooting best practices.....	161
Appendix A: Definition file templates.....	163
Setting event actions.....	163
Defining a file for running an event action.....	163
Event action definition file format.....	164
Format of the email template definition file.....	165
Command template definition files formats.....	170
Appendix B: Troubleshooting.....	176
Solving performance problems.....	176
Troubleshooting granular data collection error codes.....	180
Index.....	185

Preface

Product version

This document revision applies to Hitachi Ops Center Analyzer v10.8.2 or later.

Intended audience

This document provides an overview of the Hitachi Ops Center Analyzer software. This document is intended for storage administrators and infrastructure administrators.

Release notes

Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.com/Documents>.

Related documents

The following documents are referenced or contain more information about the features described in this manual.

- *Hitachi Ops Center Analyzer Installation and Configuration Guide*, MK-99ANA001-00
- *Hitachi Ops Center Analyzer Detail View Query Language User Guide*, MK-99ANA006-00
- *Hitachi Ops Center Analyzer Detail View Metrics Reference Guide*, MK-99ANA005-00

Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	<ul style="list-style-type: none"> Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK. Indicates emphasized words in list items.
<i>Italic</i>	<ul style="list-style-type: none"> Indicates a document title or emphasized words in text. Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: <pre>pairdisplay -g group</pre> (For exceptions to this convention for variables, see the entry for angle brackets.)
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>
< > angle brackets	Indicates variables in the following scenarios: <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> Variables in headings.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to additional information.

Icon	Label	Description
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Important	Highlights information that is essential to the completion of a task.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	CAUTION	Warns the user of a hazardous situation that, if not avoided, could result in major or minor injury.
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10 ³) bytes
1 megabyte (MB)	1,000 KB or 1,000 ² bytes
1 gigabyte (GB)	1,000 MB or 1,000 ³ bytes
1 terabyte (TB)	1,000 GB or 1,000 ⁴ bytes
1 petabyte (PB)	1,000 TB or 1,000 ⁵ bytes
1 exabyte (EB)	1,000 PB or 1,000 ⁶ bytes

Logical capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB

Logical capacity unit	Value
	Open-systems: <ul style="list-style-type: none"> ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	1,024 (2 ¹⁰) bytes
1 MB	1,024 KB or 1,024 ² bytes
1 GB	1,024 MB or 1,024 ³ bytes
1 TB	1,024 GB or 1,024 ⁴ bytes
1 PB	1,024 TB or 1,024 ⁵ bytes
1 EB	1,024 PB or 1,024 ⁶ bytes

Accessing product documentation

Product user documentation is available on the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The [Hitachi Vantara Support Website](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to the Hitachi Vantara Support Website for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: Ops Center Analyzer overview

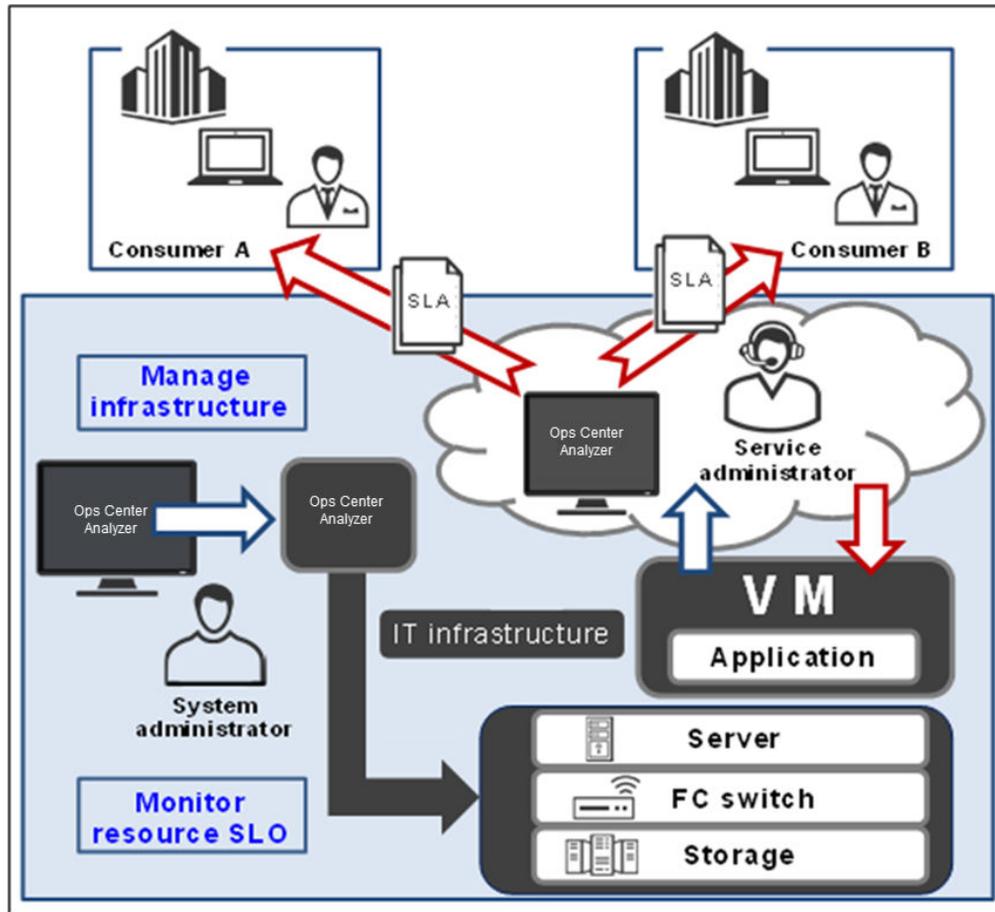
Ops Center Analyzer is data center management software for monitoring and reporting end-to-end performance from hosts through shared storage resources.

Product overview

With Hitachi Ops Center Analyzer, you can define and monitor storage service-level objectives (SLOs) for resource performance. You can identify and analyze historical performance trends to optimize storage system performance and plan for capacity growth.

Using Ops Center Analyzer, you register resources (storage systems, hosts, servers, and volumes) and set service-level thresholds. You are alerted to threshold violations and possible performance problems (bottlenecks). Using analytics tools, you find which resource has a problem and analyze its cause to help solve the problem.

The following figure shows how Ops Center Analyzer ensures the performance of your storage environment based on real-time SLOs.



The system administrator uses Ops Center Analyzer to manage and monitor the IT infrastructure based on SLOs, which match the service-implementation guidelines that are negotiated under a service-level agreement (SLA) with consumers.

Ops Center Analyzer monitors the health of the IT infrastructure using performance indicators and generates alerts when SLOs are at risk.

Having data center expertise, the service administrator uses Ops Center Analyzer to assign resources, such as VMs and storage capacity from registered storage systems, to consumer applications. This manages critical SLO violations and ensures that service performance meets the SLAs.

Features of Ops Center Analyzer

Ops Center Analyzer provides IT analytics capabilities and business benefits.

Unified infrastructure monitoring dashboard

Ops Center Analyzer dashboards are visual representations of the performance metrics of your infrastructure resources. The consolidated view allows you to quickly interpret the performance metrics and identify performance problems.

The consolidated dashboard view allows for the unified management of the server, storage, and network infrastructure resources. You can ensure the health of your data center by proactively monitoring the consumer groups, storage components, volumes, VMs, servers, and network devices. The advanced visual analytics aid in visualizing the performance data in easy-to-use graphs and charts. The visual cues allow for intuitive performance management.



The functions of the Ops Center Analyzer dashboard are as follows:

- Displays performance metrics summaries for the monitored resources.
- Displays warnings and critical alerts that need immediate action.
- Displays performance trends.
- Allows you to drill down from summary reports to detailed reports.
- Provides the ability to navigate to the E2E topology view for detailed analysis.

Advanced reporting

Ops Center Analyzer reporting capabilities enable you to monitor the infrastructure resources and assess their current performance, capacity, and utilization. Reporting data provides you the information you need to make informed business decisions and plan for future growth.

Ops Center Analyzer supports both standard and custom reporting capabilities.

In-depth reporting

Standard reports

- **Default reports:** The first time you log on to Ops Center Analyzer, the dashboard shows the following reports by default: System Status Summary for Consumers and User Resources, Event Trends, System Resource Status, and Resource Events. You can customize which reports display by default.
- **Critical reports:** Critical reports show resources in your storage infrastructure that have exceeded their thresholds. Critical reports are available for consumers, VMs, volumes, hosts, and system resources.

- **Summary reports:** Summary reports give you a high-level view of storage infrastructure resources. These reports are available for consumers, VMs, volumes, and system resources. Each summary report shows the number of resources with critical and warning alerts.
- **Capacity reports:** Capacity reports give you a measure of the capacity consumption of your data center resources and enable you to optimize the capacity usage of the existing storage resources. By monitoring the capacity reports, you can forecast future capacity requirements of your data center.
- **Other reports:** Ops Center Analyzer provides additional reports about hypervisors, switches, and system and resource events.

Custom reports

Ops Center Analyzer supports creating custom reports by running queries or by using an existing template for monitoring capacity and configuration data. You can also access the report builder from the Analyzer detail view UI to build custom reports for monitoring performance data.

In-context launch of Analyzer detail view Reports

You can access Analyzer detail view reports at any time directly from the E2E View. Simply click any resource icon and select Show Report in Analyzer detail view.

SLO management

SLOs are measurable parameters that are defined for monitoring the performance of user resources. With Ops Center Analyzer, you can evaluate, define, and customize the service-level objectives defined for the monitored resources such as volumes and VMs. By monitoring the SLOs, you can determine whether your infrastructure provides enough performance to meet the end user requirements specified in the service-level agreement.

Ops Center Analyzer lets you establish and monitor storage service-level objectives for business-critical applications and logical storage devices. When a service-level threshold is exceeded, integrated diagnostics facilitate in identifying the root cause. For storage operations, you can use the I/O control settings feature to set upper limits as a long-term solution across a range of users by consumer grade based on an SLO.

End-to-end monitoring

The E2E topology view provides detailed configuration of the infrastructure resources and lets you view the relationship between the infrastructure components. You can manually analyze the dependencies between the components in your environment and identify the resource causing performance problems. By using the topology maps, you can easily monitor and manage your resources. You can use this view to monitor resources in your data center: applications, virtual machines, servers, networks, and storage.

In the E2E view, each node represents a resource, and the connecting links represent the relationship between the infrastructure components. You can analyze a resource that is the target of analysis and all the associated resources. You can also view the alerts associated with all the related resources and trace the problem at the root level. The node-based E2E view helps you analyze the problem on the affected node and its impact on the rest of the infrastructure resources.

Problem identification and root cause analysis

Performance problems might occur because of varying system loads, applications updates, capacity upgrades, configuration changes, and inefficient management of resources in the shared infrastructure.

The Ops Center Analyzer advanced diagnostic engine aids in rapidly diagnosing, troubleshooting, and finding the root cause of performance bottlenecks.

Storage I/O controls

Storage I/O controls allow you to set and modify limits on volumes.

In the Data Center, some types of resources often require higher performance than others. For example, production servers such as database and application servers used to perform daily tasks of business organizations usually require high performance. However, if production servers experience decreased performance, productivity in business activities is negatively affected.

To prevent this from happening, the storage administrator needs to maintain the high performance of production servers. A drop in development server performance does not have as much of a negative effect on the entire organization as a drop in production server performance. In this case, you set upper limits to give higher priority to I/O activity from the production server over I/O activity from the development server to manage and control the impact of development activities.

Storage I/O controls are available in Ops Center Analyzer when Server Priority Manager is installed on your storage systems. You can invoke this function through Ops Center Automator after establishing a connection between the two servers. Alternatively, if Ops Center Automator is not installed on your storage system, you can use the CM REST API to create a script, which serves as a template that you modify for selected volumes to run the Server Priority Manager operation.

Using the I/O control setting, you can enable upper limits for the storage I/O activity of volumes that belong to consumers. The storage administrator clears the I/O control setting when the traffic between the server and storage system drops to acceptable levels. Furthermore, you have the option of limiting the data transfer rate on volumes affecting critical resources.

Set I/O control limits for the following:

- To achieve overall optimization of infrastructure resources during periods of I/O-intensive activity
- To maintain a quality-of-service benchmark for an SLO
- To prioritize I/O activity to optimize performance

About Hitachi Virtual Storage Software Block (VSSB)

The display of Virtual Storage Software Block storage information is similar to RAID, with the following exceptions:

- The Storage Node is not displayed in E2E view, but the information is available in the **Show detail** window (detail view server).
- Performance Analysis functions are not available, but the performance information is available in the **Show detail** window (detail view server)
- The Sparkline View does not support VSS Block.

VSSB resources

The volume information for VSSB storage includes:

- Compute Port
- Internode Port
- Processor
- Drive

Hitachi Ops Center Automator integration

Ops Center Analyzer supports integration with Ops Center Automator.

This support allows users to directly access the service templates in Ops Center Automator from the Execute Action window in the Ops Center Analyzer UI. When you notice a performance problem in your shared infrastructure, you can run the appropriate action or service template to resolve it.

Risk management capability

Ops Center Analyzer allows users to analyze trend projections in their infrastructure through predictive analytics risk reporting.

- **Near-term trending:** Define target metrics and predict performance over time to understand utilization trends.
- **Capacity planning:** View long-term capacity projections to plan resource allocation.



Note: The predictive analytics risk reporting capability is a licensed feature. Consult your Hitachi Vantara representative to obtain a license.

Getting started in your Ops Center Analyzer environment

Getting started tasks include logging on, accessing Analyzer detail view, and setting up your environment to monitor consumers and resources.

Logging on to Ops Center Analyzer

You can access the Ops Center Analyzer from a supported browser as follows:

- When Analyzer is registered with the Ops Center Common Services you can log in through the Ops Center portal from which you can launch Ops Center Analyzer as well as other Ops Center products using single sign-on. For information about how to configure single sign-on, see the Hitachi Ops Center Analyzer Installation and Configuration Guide.
- Log in directly using the Ops Center Analyzer URL.

Log in to Ops Center Analyzer directly as follows:

Procedure

1. Open a web browser.
2. Enter the URL for Ops Center Analyzer in the address bar:

```
http://host-where-Ops-Center-Analyzer-is-installed:port-number/Analytics/login.htm
```

where *port-number* is the port number of the Ops Center Analyzer management server. The default port number is 22015.

To access Ops Center Analyzer in secure mode, use `https`.

The default port number for secure mode is 22016.

3. Type a user ID and password.
4. Click **Log In**.

Accessing Ops Center Analyzer detail view

Use Ops Center Analyzer detail view to conduct historical trend analysis across a wide set of infrastructure statistics, create advanced monitoring custom reports, and perform additional troubleshooting and diagnostics.

You can access Analyzer detail view reports at any time directly from the E2E View. Simply click any resource icon and select Show Report in Analyzer detail view. The Analyzer detail view UI is launched in a separate browser window. The resource tree opens to the selected resource along with the latest available report in the Performance view.

You can also access Ops Center Analyzer detail view from the More Actions menu.



Note: Certain management tasks require logging directly on to the Analyzer detail view server as the `admin` user instead of using the More Actions menu (which logs on to the server as a general user). The management tasks documented in this guide state when it is necessary to log on as the `admin` user.

Use the Analyzer detail view online help to view details about reporting tasks and features.



Note: When you select one of the following resources to open in Analyzer detail view, the parent resource appears instead:

- VMware ESXi: CPU, Memory, NIC, HBA
- Hyper-V: Memory, HBA
- Storage: Others

This is because these parent and child resources are not considered independent entities in Analyzer detail view.

Launching Ops Center Automator from Ops Center Analyzer

Use Hitachi Ops Center Automator to run service templates and monitor action templates.

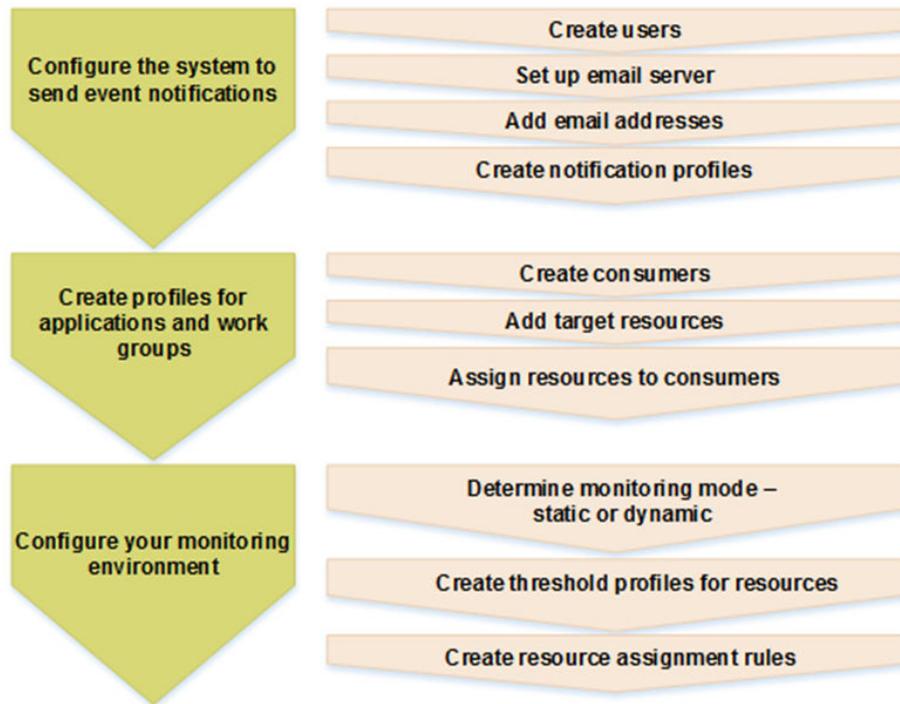
You can access Ops Center Automator when using the **Execute Action** window, typically to monitor the status of a service template you have just initiated. Selecting Launch Ops Center Automator opens the UI in a separate browser tab.

To perform administrative tasks in Ops Center Automator, such as setting permissions for Ops Center Analyzer users, launch the Ops Center Automator UI in a separate browser tab from the **Tools** menu on the main dashboard. Quick access to Ops Center Automator from the main dashboard makes the workflow convenient in that after troubleshooting performance problems in your infrastructure, you can invoke Ops Center Automator and create or edit an existing service template as a response measure to a recurring problem.

Setting up a monitoring environment for your infrastructure resources

Set up notification profiles for email alerts, create consumers to represent applications or workgroups, and establish thresholds for your infrastructure resources.

Use the following workflow to set up your monitoring environment.



- To understand consumers, resources, and thresholds, see:
 - [Overview of managing and monitoring infrastructure resources \(on page 90\)](#)
 - [Consumer settings \(on page 109\)](#)
 - [Resource monitoring settings \(on page 91\)](#)
- To set up email and notification profiles, see:
 - [Configure the email server \(on page 114\)](#)
 - [Create a condition profile \(on page 115\)](#)
- To create consumers and resources, see:
 - [Creating a consumer \(on page 110\)](#)
 - [Creating a resource assignment rule \(on page 110\)](#)
 - [Running a resource assignment rule \(on page 111\)](#)
- To monitor resources, see:
 - [Configuring resource monitoring \(on page 90\)](#)
 - [Monitoring using a static threshold \(on page 99\)](#)
 - [Monitoring using a dynamic threshold \(on page 94\)](#)
 - [Setting thresholds for user resources \(on page 103\)](#)
 - [Creating user resources assignment rule \(on page 105\)](#)
 - [Running a user resource assignment rule \(on page 105\)](#)

Chapter 2: Quick access to vital performance reports

Ops Center Analyzer communicates various types of information to you about the resources in your data center through vital performance reports.

Monitoring performance is crucial to effective data center management. System administrators require access to resource information at any moment, displayed to help understand the complex processes in the data center.

Ops Center Analyzer offers numerous gauges and reports for instant recognition of problems and long-term tracking of resource usage and events. You can also customize the display of the dashboard, charts, and reports.

Dashboard overview

The Ops Center Analyzer dashboard provides reports that display the performance status of system resources (hypervisors, storage systems, and switches), user resources (volumes, hosts, and VMs) and consumers (user resources such as virtual machines or volumes grouped by company name or business system consumer).

The dashboard shows IT infrastructure health based on real-time service-level objective (SLO) information, and provides status reports that display the capacity and performance data of all your monitored system and user resources.

When you first log on after the product is installed, the dashboard is unavailable and you are prompted for setup tasks. If you click OK, the Analyzer detail view Server window opens, and you can begin the initial setup by clicking Edit Settings.

After the initial setup, the dashboard displays the following reports by default:

- System Status Summary for Consumers and User Resources report: Displays the performance and status summary of monitored consumers, and gauges of the number of VMs, hosts, and volumes with alerts.
- Event Trends report: Displays the changes in the number of critical, warning, and information alerts for the past 72 hours.
- System Resource Status report: Displays the status of monitored server, SAN, and storage components.
- Resource Events report: Displays a list of resource events based on the time of event occurrence. The most recent events appear at the top of the list.

You can customize the dashboard to display reports that you prefer to monitor. To reset the dashboard to display the default reports, click Restore Default Settings.

The status indicator displays the number and severity of alerts generated by the storage environment.

Icon	Color	Definition
	Red	Critical
	Yellow	Warning
	Blue	Information
	Green	Normal

Overall performance summary

If the summary bar is green, then all SLOs have been met. However, these status and empty reports can also indicate that you did not set up the Ops Center Analyzer for performance monitoring. Use the configuration workflow to set up the Ops Center Analyzer to monitor managed resources.

If the summary bar is red, a critical error was detected. Alert bars are listed in order of priority:

- Consumer
- User Resource
- System Resource

The dashboard displays the critical and warning status in the system in order of priority.

For example, if a warning alert occurs in a Consumer or User Resource, and a critical alert occurs in a System Resource, the indicator bar displays the status of the Consumer or User Resources because Consumers and User Resources have a higher priority than System Resources.

The default refresh time for the dashboard is 5 minutes. You can manually refresh the dashboard by clicking **Refresh** at the top of the dashboard, or by configuring the refresh interval time in the **Dashboard Settings** window.



Note: If a VM on a Hyper-V server is set to be monitored, the relationships between hosts and storage systems do not display on the E2E View.

Print

Click the Print icon to save an output of the current dashboard to a file.



Note: The Chrome browser will save the output as a PDF file. Internet Explorer and Firefox only support the IPS file format. If you have a suitable PDF software package installed, you can use the native print function of Internet Explorer or Firefox and select to print as a PDF file.

Search

The search feature on the home page lets you search for a resource in the Consumers, Servers, Storage Systems, and Volumes categories. From the returned search results, select the resources you need to analyze, and launch the E2E view or Sparkline view for further analysis.

System Status Summary for Consumers and User Resources

The System Status Summary for Consumers and User Resources report displays the performance status summary of monitored consumers, VMs, hosts, and volumes.



System Status Summary for Consumers

Consumers are business management units under which user resources such as virtual machines, hosts, or volumes can be grouped by company name or business system consumer, and assigned grades based on their importance.



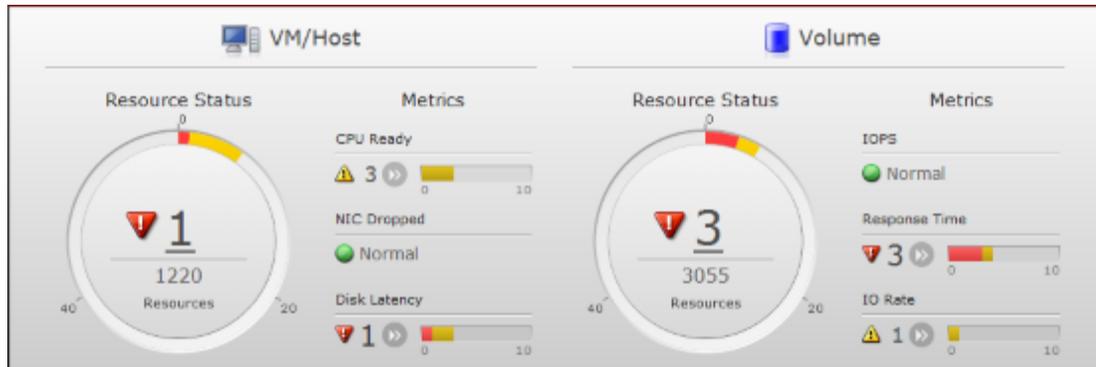
The Consumer pane displays the consumer grades, the total number of alerts for each, and a bar graph of the highest alert severity in each grade, as follows:

- Consumer grades are listed from the highest to the lowest grade. Consumers with a bronze grade are grouped under Others.
- The total number of critical or warning alerts displays above the total number of consumers for each consumer grade. If there are critical and warning alerts for a specific consumer grade, then the alert for the highest severity displays.

- The number of alerts for each grade is a link that opens the **Consumers - Critical/Warning** window.
- If more than one alert exists for any specific grade, a bar graph appears. A red bar indicates critical errors, and a yellow bar is a warning. If the number of resources with critical and warning alerts exceeds 10, only the critical alerts display.
- If there are no consumers associated with a specific type of grade, - / - displays for the number of alerts.
- If no consumers are associated with a grade, or if no data is available, **No Data** displays instead of a bar graph.

System Status Summary for User Resources

The System Status Summary - VM/Host and Volume reports display the status of all monitored VMs, hosts, and volumes.



Both the VM/Host pane and the Volume pane display a Resource Status information gauge, where the top number is the total critical or warning alerts received from the VMs and hosts, or volumes that exceeded the critical or warning thresholds for any monitored metric. The bottom number indicates the total number of VMs, hosts, or volumes in the system.

Under Metrics, a bar graph displays the total number of VMs and hosts with critical and warning alerts for any monitored metric.

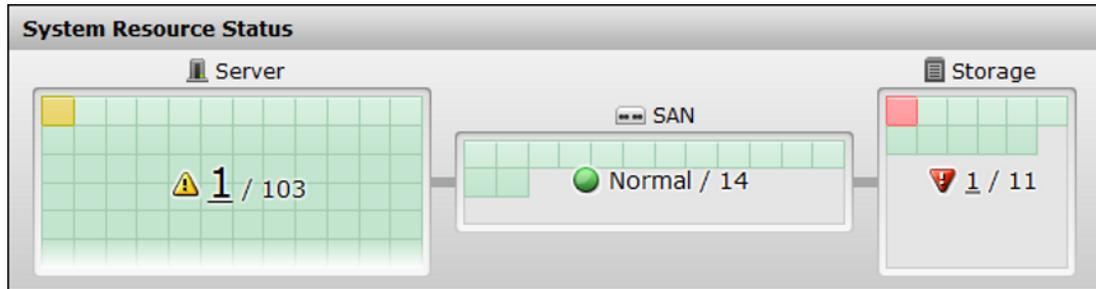
For the Volume summary, Metrics displays a bar graph of the total number of volumes with critical and warning alerts for any metrics. In the bar graph, red is the total number of critical alerts and yellow is the number of warning alerts.

For example, if there are 3 critical alerts and 5 warning alerts, then both critical and warning display in the gauge. If the number of critical alerts is greater than 9, only the red bar (critical) displays because the maximum value of the gauge is 10.

To view details about the resources that exceed the defined critical or warning thresholds, click the number link in the information gauge chart or bar graph. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration to analyze the performance problem.

System Resource Status report

The System Resource Status report is one of the default reports that appears on the Ops Center Analyzer dashboard. It provides a heat map of the current status of system resources such as server (CPU, memory, NIC, HBA, and disk), SAN (switches), and storage (ports, processors, cache, pools, and parity groups) components.



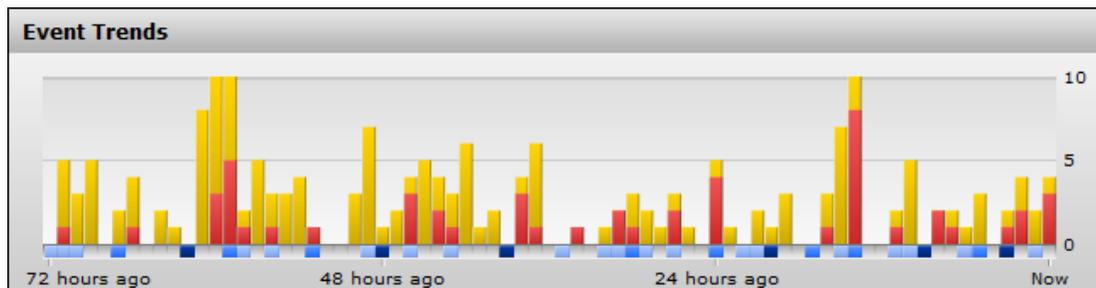
Each red tile shows a critical alert, and each yellow tile shows a warning alert. Unknown resources are considered Normal and are represented by green tiles.

To view details about the resources that exceeded the defined thresholds, click the number link. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Event Trends report

The Event Trends report is one of the default reports that appears on the Ops Center Analyzer dashboard. It provides a trend report of all critical and warning alerts in your environment for the past 72 hours. It is useful in comparing the change history and the number of critical alerts, especially for an administrator who manages the entire system.

- When you hover over a specific bar in the trend report, a tool tip displays the total number of alerts, the number of critical alerts, and the number of warning alerts for a specific time period. Each bar represents an hour.
- Blue blocks at the bottom of the trend report indicate configuration changes for a specific time period. The darker the shade of blue, the larger the configuration change from one time period to the next.
- When you hover over a change history at the bottom of the trend report, the tool tip displays the number of changes in a specific time period. The date is in the format yyyy-mm-dd, and the time hh:mm.



If you see a spike on the trend report, navigate to the **E2E View** from the System Status Summary report to view the data center topology and review the configuration and status information. You can then use the Sparkline View and analytics workflow to solve the problem.

Resource Events report

The Resource Events report table on the Ops Center Analyzer dashboard lists the most current resource events in descending order, based on the Date Time column. You can view a maximum of 500 critical and warning events that occurred in the past 24 hours.

Resource Events					
Column Settings					
Level	Message	Date Time	Category	Device Name	Component Name
Warning	"Virtual Nic Dropped IP Packets..."	2016-04-11 01:22:00	Performance	172.19.36.43	HCSBeta
Critical	"Virtual Nic Dropped IP Packets..."	2016-04-11 01:07:00	Performance	172.19.36.43	HCSBeta
Warning	"Virtual Nic Dropped IP Packets..."	2016-04-11 00:37:00	Performance	172.19.36.43	HCSBeta
Critical	"Virtual Nic Dropped IP Packets..."	2016-04-11 00:07:00	Performance	172.19.36.43	HCSBeta
Critical	"CPU Ready" changed from "Wa..."	2016-04-10 23:59:00	Performance	172.19.36.62	172.19.36.62

To view details about the resources that exceeded the defined thresholds, click the associated message link. Click Show E2E View from the Event Detail tab to view the data center topology and review the configuration and status information.

System and Resource Events

You can view the latest events in one place and manage the events based on the status.

The Events tab displays details about significant events in your monitored environment.

The screenshot shows the 'Events' tab in the Hitachi Ops Center Analyzer. The interface includes a navigation bar with 'Dashboard', 'Analytics', 'Operations', 'Events', and 'Administration'. Below the navigation bar, there are tabs for 'All Events', 'Resource Events', and 'System Events'. The main content area displays a table of events with columns for Level, Message, Date Time, Category, Device Name, and Component Name. The table is filtered to show 100 rows out of 1,360. The events listed include various performance and critical alerts, such as 'Dropped Rx (VMware Virtual Machine)', 'Network Usage (VMware Virtual Machine)', and 'Write Pending Rate (MPB CLPB)'.

Level	Message	Date Time	Category	Device Name	Component Name
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Normal to W...	2017-03-22 17:05:00	Performance	172.19.36.60	Hitachi Storage Provider fo...
Information	"Network Usage (VMware Virtual Machine)" changed from Critical to ...	2017-03-22 17:05:00	Performance	172.19.36.60	HIAAhdca Server 7
Critical	"Dropped Rx (VMware Virtual Machine)" changed from Warning to C...	2017-03-22 17:02:00	Performance	172.19.36.61	Win_Ramju
Critical	"Dropped Rx (VMware Virtual Machine)" changed from Warning to C...	2017-03-22 17:02:00	Performance	172.19.36.61	HSV-Perf-Collector
Critical	"Dropped Rx (VMware Virtual Machine)" changed from Warning to C...	2017-03-22 17:02:00	Performance	172.19.36.61	HSV
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Warning to C...	2017-03-22 17:02:00	Performance	172.19.36.63	PentahDenver
Information	"Dropped Rx (VMware Virtual Machine)" changed from Warning to N...	2017-03-22 17:00:00	Performance	172.19.36.60	HDCAv7ARMaster
Information	"Write Pending Rate (MPB CLPB)" changed from Critical to Normal in...	2017-03-22 17:00:00	Performance	VSP G1000 (S6006)	0
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Critical to Wa...	2017-03-22 17:00:00	Performance	172.19.36.61	Win_Ramju
Information	"Dropped Rx (VMware Virtual Machine)" changed from Warning to N...	2017-03-22 17:00:00	Performance	172.19.36.60	HIAAhdca Server 7
Information	"Dropped Rx (VMware Virtual Machine)" changed from Warning to N...	2017-03-22 17:00:00	Performance	172.19.36.60	Hitachi Storage Provider fo...
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Critical to Wa...	2017-03-22 17:00:00	Performance	172.19.36.61	HSV
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Critical to Wa...	2017-03-22 17:00:00	Performance	172.19.36.61	HSV-Perf-Collector
Critical	"Network Usage (VMware Virtual Machine)" changed from Normal to W...	2017-03-22 17:00:00	Performance	172.19.36.60	HIAAhdca Server 7
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Critical to Wa...	2017-03-22 17:00:00	Performance	172.19.36.63	PentahDenver
Information	"Write Pending Rate (MPB CLPB)" changed from Warning to Normal...	2017-03-22 16:58:00	Performance	VSP G1000 (S6006)	0
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Normal to W...	2017-03-22 16:57:00	Performance	172.19.36.60	HIAAhdca Server 7
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Normal to W...	2017-03-22 16:57:00	Performance	172.19.36.60	Hitachi Storage Provider fo...
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Normal to W...	2017-03-22 16:57:00	Performance	172.19.36.60	HDCAv7ARMaster
Information	"Dropped Rx (VMware Virtual Machine)" changed from Warning to N...	2017-03-22 16:56:00	Performance	172.19.36.60	Hitachi Storage Provider fo...
Information	"Dropped Rx (VMware Virtual Machine)" changed from Warning to N...	2017-03-22 16:56:00	Performance	172.19.36.60	HIAAhdca Server 7
Information	"Dropped Rx (VMware Virtual Machine)" changed from Warning to N...	2017-03-22 16:56:00	Performance	172.19.36.60	HDCAv7ARMaster
Critical	"Write Pending Rate (MPB CLPB)" changed from Warning to Critical...	2017-03-22 16:54:00	Performance	VSP G1000 (S6006)	0
Warning	"Write Pending Rate (MPB CLPB)" changed from Normal to Warning...	2017-03-22 16:52:00	Performance	VSP G1000 (S6006)	0
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Normal to W...	2017-03-22 16:50:00	Performance	172.19.36.60	HIAAhdca Server 7
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Normal to W...	2017-03-22 16:50:00	Performance	172.19.36.60	Hitachi Storage Provider fo...
Warning	"Dropped Rx (VMware Virtual Machine)" changed from Normal to W...	2017-03-22 16:50:00	Performance	172.19.36.60	HDCAv7ARMaster
Warning	"Write Pending Rate (MPB CLPB)" changed from Normal to Warning...	2017-03-22 16:50:00	Performance	VSP G1000 (S6006)	0

There are two categories of events:

- **System Events**

The System Events tab displays Management and Event Action events generated when system settings must be verified or configured.

- **Resource Events**

The Resource Events tab displays Performance events generated when a device or component (server, storage system, network device, and so on) does not perform optimally.

You can analyze the Resource events by using the end-to-end network topology view to identify the resource that generated the event.

The All Events tab displays both System and Resource events. Each event indicates the level of the alert, the date and time of the alert message, category, device name, and component name. Click a message in the Message column to open the **Event Detail** window.

Use the Event Detail window to display more event details, such as the device type and component type. You can scroll through the list for more events. For Resource events, you can click Show E2E View to view the network topology.

The Event levels classifications are as follows:

- **Critical:** Event that requires immediate attention
- **Warning:** Event that might become critical in the future
- **Informational:** No immediate action required

Sending dashboard reports to users

You can schedule the delivery of dashboard reports to users. All current dashboard reports are sent as an email attachment.

Before you begin

Make sure the mail server and the sender address are set up as described in [Configuring the mail server \(on page 114\)](#).

Procedure

1. Click **Dashboard Settings**.
2. Set **Send Reports** to **ON**.
3. Choose to deliver the report **Daily**, **Weekly** (Sun-Sat), or **Monthly**, and set the **Time** of day when they are to be sent.
4. Enter one or more email recipients in the **To** field.
5. Click **OK**.

Customizing the dashboard

Customize the dashboard settings to display the reports you prefer to see when you log in. You can also schedule the email distribution of reports.

User permissions for managing dashboard settings

The dashboard settings that you can configure depend on the user permissions assigned to you.

- **System account user**

A system account is a fully privileged built-in account used to log on after the product is installed. After the initial setup, the dashboard by default displays built-in reports. A system account user can make changes to the default dashboard and manage the dashboard organization for all user accounts. The changes made to the default dashboard is reflected on the dashboards of other users using the default dashboard.

- **Admin or modify users**

When users with admin or modify permissions log on for the first time, they see the default dashboard configured by the system account user. The users with admin or modify permissions cannot edit the default dashboard settings, but can create their own custom dashboard.

Configuring dashboard settings

Configure the following settings using the Dashboard Settings window:

- **Toggle between default and custom dashboards**

Users with admin or modify permissions can toggle between default and custom dashboards.

- **Set refresh interval**

To set the data refresh interval, enter the time period in the Refresh Interval text box. The default refresh interval is 5 minutes. The reports refresh automatically at the set interval when the autorefresh option is enabled. The built-in reports are enabled for autorefresh by default. You can either enable or disable the autorefresh function for custom reports. Check whether reports are enabled for autorefresh in the Refresh column.

- **Restore default settings**

To reset the dashboard to display the default settings, click Restore Default Settings. (This will not affect Send Reports settings.)

- **Schedule delivery of reports**

To deliver a copy of the dashboard reports on a scheduled basis to users, set Send Reports to ON. You can then choose the date and time and one or more recipients.

- **Display reports**

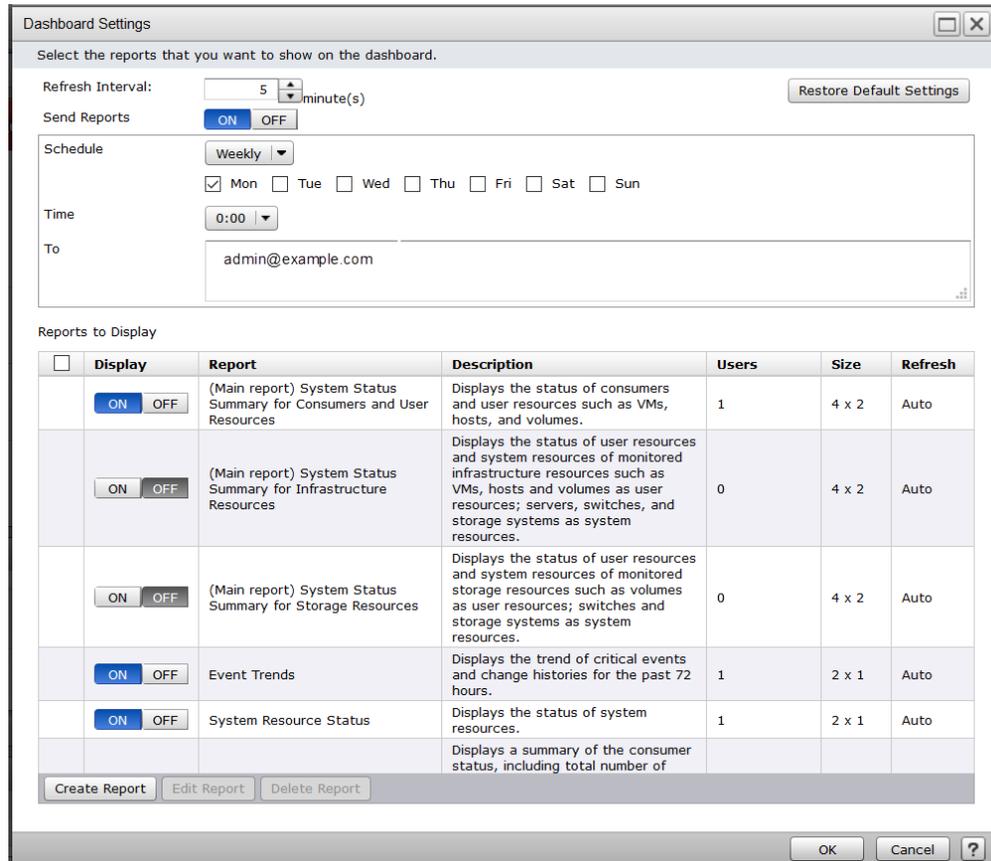
To display the reports you want to view on your dashboard, set the ON or OFF options.

▪ **Create and manage custom reports**

Create custom reports for monitoring capacity and configuration data by clicking the Create Report option. Only the system account user can create, edit, and delete public reports that are accessible by all users. Users with admin or modify permissions can create private reports for display on their dashboard.

▪ **Notification of new reports**

When a system account user creates a new custom report, it is made available for all users. The New indicator appears next to the report name.



Chapter 3: Creating custom reports

Ops Center Analyzer enables you to create custom reports for monitoring capacity and configuration data. Select the report metrics that you want to monitor, pick the report format, and set the report definitions for displaying your data.

Custom reports overview

Ops Center Analyzer allows you to create your own reports for monitoring capacity and configuration data.

Public and Private reports

Ops Center Analyzer supports creating public and private report types:

- Built-in reports: Ops Center Analyzer provides a collection of default reports that are referred to as built-in reports. Built-in reports cannot be edited.
- Public reports: Public reports can be accessed by all users. Only the system account user can create, edit, or delete public reports. Other users have only view access.
- Private reports: Private reports are accessible only by the users who create them. Users with admin or modify permissions can create, edit, and delete private reports.

User permissions for creating custom reports

Ops Center Analyzer user roles and permissions for creating custom reports are as follows:

- System: A system account user can create, edit, and delete public reports.
- Admin or Modify: A user with admin or modify permissions can create, edit, and delete custom reports created for private access. In addition to the default reports, these users can also view the public reports created by the system account user.

Custom report creation methods

Create a custom report using any of following methods:

- By entering a query: Write a query to create a report. For information about how to write a query, see the *Hitachi Ops Center Analyzer Detail View Query Language User Guide*. For information about capacity and configuration attributes, see the *Hitachi Ops Center Analyzer Detail View Metrics Reference Guide*.
- By using a report template: Select a report template and customize the report definitions to create a new report. A system account user can access the private report templates created by users with admin or modify permissions. In addition to the built-in report templates, users with admin or modify permissions can access the public report templates created by the system account user.

Custom report format

View your report data in the following formats:

- Line graphs display the data trends over a selected time period. Create combination graphs using the primary and secondary y-axis and a common x-axis. Use the combination graphs to display data for one or more data series.
- Table graphs display data in tabular format for the selected metrics over a selected time period.

Custom report preview

Preview the report before publishing and make necessary modifications.

Custom report notification

When a system account user creates a new custom report, other users are notified by the following methods:

- On the Dashboard tab, the New status indicator appears next to the Dashboard Settings button.
- In the Dashboard Settings window, under Reports, you can view the newly added reports marked as New.

Creating custom reports

Create custom reports for monitoring capacity and configuration data.

Before you begin

- Only a system account user can create public reports that are accessible by all users. Public reports can only be edited or deleted by a system account user. In addition to the built-in report templates, a system account user can access the private report templates created by users with admin or modify permissions.
- Users with admin or modify permissions can only create reports for private access. Private reports can be edited or deleted by users who create them. In addition to the built-in report templates, users with admin or modify permissions can access the public report templates created by the system account user.

Procedure

1. From the **Dashboard** tab, click **Dashboard Settings**.
2. In the **Dashboard Settings** window, click **Create Report**.
3. In the **Create Report** window, either enter a query or select a report template for creating reports.
4. To create a custom report by selecting a report template:
 - a. Browse the report templates from the **Copy From** list. When you hover over the report names in the **Copy From** list, the details such as report name, description, size, query, and target time period of analysis are displayed to help you select an appropriate template.
 - b. Select a report template from the **Copy From** list. The report properties are autofilled with predefined values from the selected template.
 - c. Edit the predefined report properties such as report name, description, size, query, and other report properties.
5. To create a custom report by entering a query:
 - a. Enter the report name and description.
 - b. Select the report type and size. View the report results using **Line Graphs** or **Table**.
 - c. In the Query text box, enter a query.
 - For information about how to write queries, see the *Hitachi Ops Center Analyzer Detail View Query Language User Guide*.
 - For information about Hitachi Enterprise Storage system capacity and configuration metrics and attributes, see the *Hitachi Ops Center Analyzer Detail View Metrics Reference Guide*.
 - For information about all other resource attributes, go to the Analyzer detail view UI, and from the application bar, click the **Manage** icon. Under Administration, click **Show Schema**. The **Show Schema** window displays the details of all metrics, resource definitions, and relationships.
 - d. Select the **Time Period of Analysis**.
6. Select **Auto Refresh** to enable autorefresh for the report. Reports refresh automatically at the set refresh interval time. If you disable this option, you must manually refresh the report by using the refresh icon at the top of the dashboard, or the refresh icon at the top of each report.

7. To view the query results, click **Set Report Definitions**. Set the report definitions for displaying your data either using line graph or table format.



Note: For tabular data, you can use the **Default Sort** column to set the sorting behavior for each metric (ascending or descending) in dashboard reports that are sent to users.

8. To review the report fields, report design, and other details before publishing, click **Preview**.
9. Click **OK**.
The report appears on the dashboard.

Managing custom reports

Edit or delete the custom reports, depending on your access permissions.

Before you begin

- Only a system account user can create public reports that are accessible by all users. Public reports can only be edited or deleted by a system account user.
- Users with admin or modify permissions can only create reports for private access. Private reports can be edited or deleted by the users who create them.

Procedure

1. From the Dashboard tab, click **Dashboard Settings**.
2. In the **Dashboard Settings** window, you can create, edit, and delete custom reports.
 - To edit a custom report, select the report, and then click **Edit Report**. Modify the report definitions such as report name, description, query, size, type, time period, and other details.
 - To delete custom reports, select the reports, and then click **Delete Report**.

If you are a system account user, before deleting the public reports, check how many users are viewing this report on their custom dashboard by verifying the number of users in the **Users** column. Hover over the number of users to see details about users viewing this report. Make sure you notify the users before deleting custom reports.

Examples of creating custom reports

Ops Center Analyzer enables you to write queries and build your own reports. Use the query statements to retrieve information from the Analyzer detail view database and create configuration and capacity reports.

Creating a line graph

Plot a line graph to visualize the data patterns over time. You can plot line graphs for one or more data series.

Before you begin

- Only a system account user can create public reports that are accessible by all users. Public reports can only be edited or deleted by a system account user.
- Users with admin or modify permissions can only create reports for private access. Private reports can be edited or deleted by the users who create them.

This example shows how to create a custom report by using the report template. You will create a line graph report for monitoring the used capacity and free capacity of Hitachi Dynamic Tiering (HDT) pools for VSP G series storage systems.

Procedure

1. From the Dashboard tab, click **Dashboard Settings**.
2. In the **Dashboard Settings** window, click **Create Report**.
3. In the **Create Report** window, either enter a query or select a report template for creating reports.
4. Browse through the report templates listed in the **Copy From** list. Hover over the report names in the list to view details such as report name, description, size, query, target time period, and other details which help you select an appropriate template.
5. From **Copy From** list, select [Line]Capacity of Drive Types by Storage System (HDT).

The report attributes such as name, description, type, size, query statement, and time period of analysis are autofilled with predefined values from the selected template. The query statement associated with the report template is autofilled in the **Query** text box. Edit the predefined values to customize your report.

Create Report

Select the template of an existing report or enter a query to create a new configuration report.

Copy From: [Built-In] Capacity of Drive Types by Storage System (HDT) |

Name: * Capacity of Drive Types by Storage System (HDT)

Description: Displays the capacity information of drive types by storage system. This report shows the capacity of HDT pools.

Type:

 Table

 Line Graph

Size: 4 x 1

Query: ? * raidStorage [=modelName rx .*] [=serialNumber rx .*] [@All_ThinFree rx b .*] {inputCounters=raidPoolTier.tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB} [@All_ThinUsed rx b .*] {inputCounters=raidPoolTier.tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}

Time Period: Last 24 hours

Auto Refresh

Set Report Definitions

* Required

OK Cancel ?

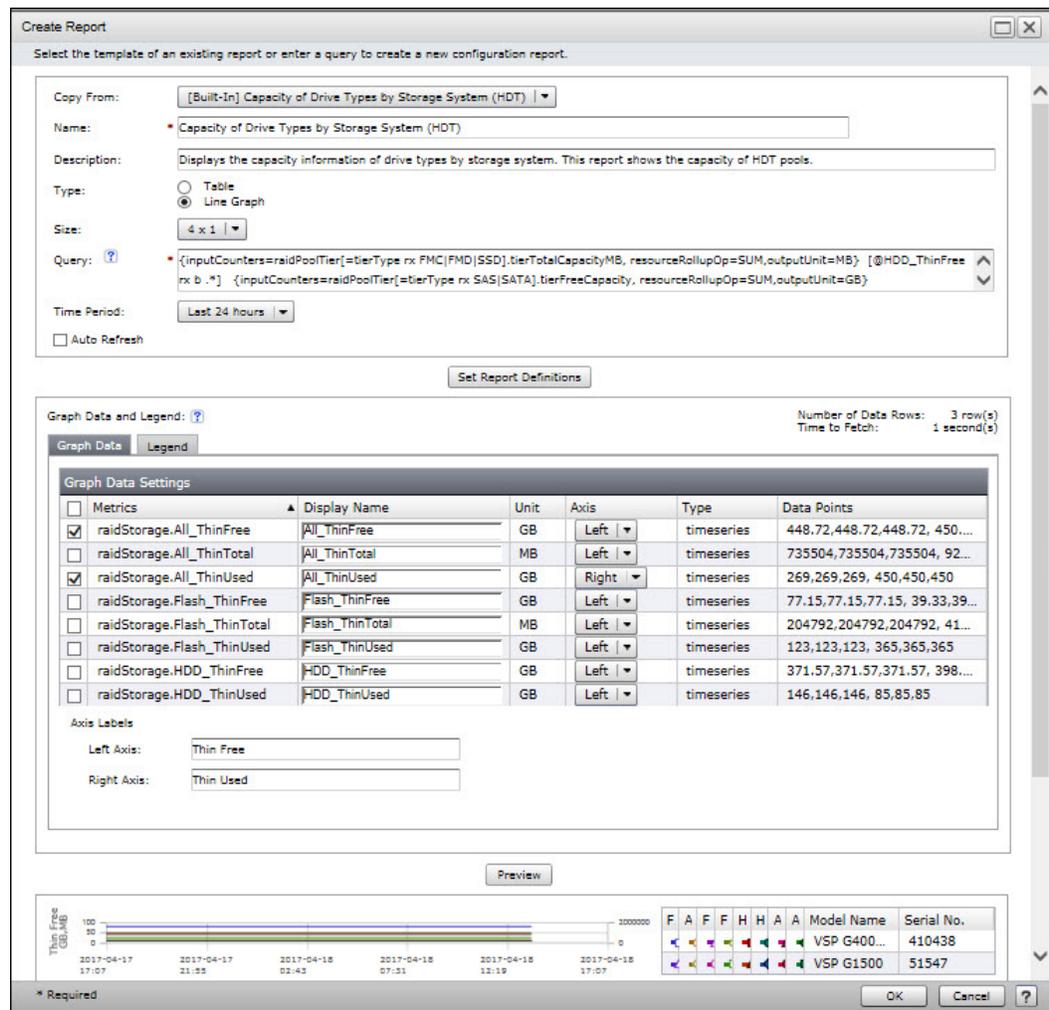
The sample query for creating a line graph for monitoring the capacity of HDT pools of raidstorage systems is as follows:

```
raidStorage [=modelName rx .*] [=serialNumber rx .*] [@All_ThinFree
rx b .*] {inputCounters=raidPoolTier.tierFreeCapacity,
resourceRollupOp=SUM,outputUnit=GB} [@All_ThinUsed rx b .*]
{inputCounters=raidPoolTier.tierUsedCapacity, resourceRollupOp=SUM,
outputUnit=GB} [@All_ThinTotal rx b .*]
{inputCounters=raidPoolTier.tierTotalCapacityMB, resourceRollupOp=SUM,
outputUnit=MB} [@Flash_ThinFree rx b .*]
{inputCounters=raidPoolTier[=tierType rx FMC|FMD|
SSD].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@Flash_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx FMC|
FMD|SSD].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@Flash_ThinTotal rx b .*]{inputCounters=raidPoolTier[=tierType rx FMC|
FMD|SSD].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
[@HDD_ThinFree rx b .*] {inputCounters=raidPoolTier[=tierType rx SAS|
SATA].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@HDD_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx SAS|
SATA].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@HDD_ThinTotal rx b .*] {inputCounters=raidPoolTier[=tierType rx SAS|
SATA].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
[@FMC_ThinFree rx b .*] {inputCounters=raidPoolTier[=tierType rx
FMC].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@FMC_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx
FMC].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@FMC_ThinTotal rx b .*] {inputCounters=raidPoolTier[=tierType rx
FMC].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
[@FMD_ThinFree rx b .*] {inputCounters=raidPoolTier[=tierType rx
FMD].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@FMD_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx
FMD].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@FMD_ThinTotal rx b .*] {inputCounters=raidPoolTier[=tierType rx
FMD].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
[@SSD_ThinFree rx b .*] {inputCounters=raidPoolTier[=tierType rx
SSD].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@SSD_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx
SSD].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@SSD_ThinTotal rx b .*] {inputCounters=raidPoolTier[=tierType rx
SSD].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
[@SAS_ThinFree rx b .*] {inputCounters=raidPoolTier[=tierType rx
SAS].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@SAS_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx
SAS].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@SAS_ThinTotal rx b .*] {inputCounters=raidPoolTier[=tierType rx
SAS].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
[@SATA_ThinFree rx b .*] {inputCounters=raidPoolTier[=tierType rx
SATA].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@SATA_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx
SATA].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@SATA_ThinTotal rx b .*] {inputCounters=raidPoolTier[=tierType rx
```

```
SATA].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
[@External_ThinFree rx b .*] {inputCounters=raidPoolTier[=tierType rx
External].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@External_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx
External].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@External_ThinTotal rx b .*] {inputCounters=raidPoolTier[=tierType rx
External].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
[@Mixed_ThinFree rx b .*] {inputCounters=raidPoolTier[=tierType rx
Mixed].tierFreeCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@Mixed_ThinUsed rx b .*] {inputCounters=raidPoolTier[=tierType rx
Mixed].tierUsedCapacity, resourceRollupOp=SUM,outputUnit=GB}
[@Mixed_ThinTotal rx b .*] {inputCounters=raidPoolTier[=tierType rx
Mixed].tierTotalCapacityMB, resourceRollupOp=SUM,outputUnit=MB}
```

6. Click **Set Report Definitions** to add graph data and legend entries.

The query results display the key metrics for monitoring the enterprise-wide HDT total pool capacity, used capacity, and free capacity details.



7. In the **Graph Data** tab, select the key metrics to represent on the left and right axes.

- a. Select the performance metric for the left-axis. Select **raidstorage.All_ThinUsed** to represent used capacity of the HDT pool on the left-axis. Specify a display name, and from the **Axis** list, select **Left**.
 - b. Select the performance metric for the right-axis. Select **raidstorage.All_ThinFree** to represent free capacity of the HDT pool on the right-axis. Specify a display name, and from the **Axis** list, select **Right**.
8. Enter the axis labels:
 - Left Axis: `Thin Free`
 - Right Axis: `Thin Used`
 9. In the **Legend** tab, select the legend entries to appear in the graph.
 10. To review the graph data, layout, legend entries, and other details before publishing, click **Preview**.
 11. To preview the original size of the graph, click **Size Preview**.



12. Click **OK**.
The new report appears on your dashboard.

Creating a table report

Create custom reports in table format to display configuration and capacity data.

This example shows how to create a table report by using the report template. You will create a table report for monitoring the capacity consumption of the monitored storage systems in your data center.

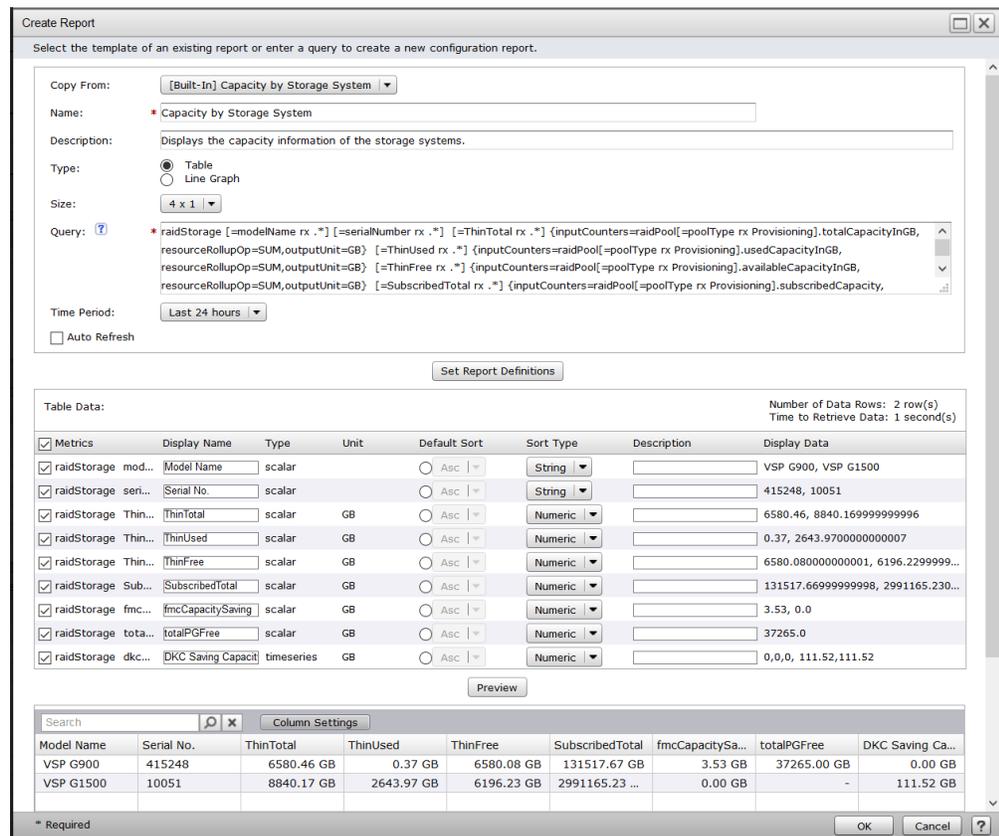
Before you begin

- Only a system account user can create public reports that are accessible by all users. Public reports can only be edited or deleted by a system account user.
- Users with admin or modify permissions can only create reports for private access. Private reports can be edited or deleted by the users who create them.

Procedure

1. From the Dashboard tab, click **Dashboard Settings**.
2. In the **Dashboard Settings** window, click **Create Report**.
3. In the **Create Report** window, either enter a query or select a report template for creating reports.
4. Browse through the report templates listed in the **Copy From** list. Hover over the report names in the list to view details such as report name, description, size, query, and target time period to help you select an appropriate template.
5. From **Copy From** list, select `[Table]Capacity by Storage System`.

The report attributes such as name, description, type, size, query statement, and time period of analysis are autofilled with predefined values from the selected template. The query statement associated with the report template is autofilled in the **Query** text box. Edit the predefined values to customize your report.



The sample query for creating a table report for monitoring capacity by storage system is as follows:

```
raidStorage [=modelName rx .*] [=serialNumber rx .*] [=ThinTotal rx .*] {inputCounters=raidPool[=poolType rx Provisioning].totalCapacityInGB, resourceRollupOp=SUM,outputUnit=GB} [=ThinUsed rx .*] {inputCounters=raidPool[=poolType rx Provisioning].usedCapacityInGB, resourceRollupOp=SUM,outputUnit=GB} [=ThinFree rx .*] {inputCounters=raidPool[=poolType rx Provisioning].availableCapacityInGB, resourceRollupOp=SUM, outputUnit=GB} [=SubscribedTotal rx .*] {inputCounters=raidPool[=poolType rx Provisioning].subscribedCapacity, resourceRollupOp=SUM,outputUnit=GB} [=fmcCapacitySaving rx .*] {inputCounters=raidPool[=poolType rx Provisioning].fmcPoolCapacitySaving, resourceRollupOp=SUM, outputUnit=GB} [@dkcSavingCapacity rx b .*] {inputCounters=raidPool[=poolType rx Provisioning].dkcSavingCapacity, resourceRollupOp=SUM,outputUnit=GB} [=totalPGFree rx .*] {inputCounters=raidPG.totalFreeSpaceInGB, resourceRollupOp=SUM, outputUnit=GB}
```

6. Click **Set Report Definitions** to add table data.

The query results display the key metrics for monitoring the capacity consumption of all monitored storage systems.

7. In the **Table Data** tab, select the key performance metrics that you want to monitor.

8. To review the table report details before publishing, click **Preview**.

9. Click **OK**.

The new report appears on your dashboard.

Model Name	Serial No.	ThinTotal	ThinUsed	ThinFree	SubscribedTotal	fmcCapacitySaving	totalPGFree	DKC Saving Capacity
VSP G400/G600	410438	2712 GB	410 GB	2302 GB	140600 GB	50 GB	105007 GB	245.08 GB
VSP G1500	51547	6349 GB	924 GB	5425 GB	450913 GB	0 GB	-	508.37 GB
VSP G1500	10051	18019 GB	6619 GB	11399 GB	194651 GB	0 GB	-	-

Sample queries

Example queries for creating custom reports for monitoring capacity and configuration data are as follows:

To monitor unallocated capacity:

```
raidStorage[=name rx .*]&[=unAllocatedCapacity rx .+]  
{resourceRollupOp=sum,inputCounters=raidStorage.unAllocatedCapacityInGB}
```

To monitor capacity in GB:

```
raidStorage[=name rx .*]&[=capacity rx .+]{resourceRollupOp=sum,  
inputCounters=raidStorage.capacityInGB}
```

To show capacity data in GB and monitor unallocated capacity:

```
raidStorage[=name rx .*]&[=capacity rx .+]{resourceRollupOp=sum,  
inputCounters=raidStorage.capacityInGB}&[=unAllocatedCapacity rx .+]  
{resourceRollupOp=sum,inputCounters=raidStorage.unAllocatedCapacityInGB}
```

To monitor the storage used capacity, unallocated capacity, allocated capacity, free capacity, and reserved capacity:

```
raidStorage[=name rx .*]&{resourceRollupOp=sum,  
inputCounters=raidStorage.capacityInGB}&{resourceRollupOp=sum,  
inputCounters=raidStorage.unAllocatedCapacityInGB} [=capacityInGB  
rx .+]&[=allocatedCapacityInGB rx .+]&[=unAllocatedCapacityInGB  
rx .+]&[=totalFreeSpaceInGB rx .+]&[=openReservedCapacityGB  
rx .+]&[=otherCapacityGB rx .+]
```

To monitor the total capacity, used capacity, and available capacity for all pools:

```
raidStorage/raidPool[=name rx .*]&[=availableCapacityInGB  
rx .+]&[=totalCapacityInGB rx .+]&[=usedCapacityInGB rx .+]
```

Chapter 4: Analyzing performance problems

Ops Center Analyzer provides analytical diagnostics to quickly identify, isolate, and determine the root cause of problems.

The traditional approach of troubleshooting performance problems in the unified infrastructure poses several challenges. For example, it can be difficult to identify performance problem in a storage infrastructure environment that includes various virtual machines, servers, network, and storage.

Ops Center Analyzer offers an out-of-the-box analytics solution which lets you identify and troubleshoot performance problems at the node level. The topology view shows the graphical representation of the infrastructure components and their dependencies. This view is crucial for troubleshooting performance problems and helps achieve efficient root cause analysis.

Identifying performance problems

The IT infrastructure is becoming more complex each day with rapidly emerging converged infrastructures. Performance problems occur due to various factors in your environment. Identifying the performance problems and troubleshooting the problems quickly is crucial.

As part of your performance management strategy, you define performance goals and criteria for monitoring your environment. The performance problems occur when these predefined goals are not met. Use Ops Center Analyzer advanced troubleshooting features to quickly fix problems.

The following situations indicate a performance problem in your environment:

- An SLO violation occurs
Typically, SLAs define the SLOs to evaluate the quality of service. SLO profiles define the threshold values for the performance parameters that you use to evaluate the quality of service. When the threshold values are exceeded, an SLO violation occurs.
- A sharp deviation from the baseline data occurs
When no SLOs are defined for your environment, you can use the baseline values to evaluate your system performance. The current performance is compared to the past performance trends, and when there is a significant deviation from the baseline values, Ops Center Analyzer sends an alert to notify you of a potential performance problem so you have enough time to troubleshoot.
- The customer notifies you of an application performance degradation and slowdown of the infrastructure.

The common causes for performance problems are as follows:

- Increased load in an otherwise stable operating environment
- Inefficient load balancing strategy, which might cause underutilization of resources
- Changes in the system configuration
- Resource management in a shared infrastructure

Infrastructure components and key performance metrics

You must analyze the key performance metrics relevant to the problem and the workload being analyzed.

The components and key performance metrics available in Ops Center Analyzer for monitoring performance are listed in the following tables:

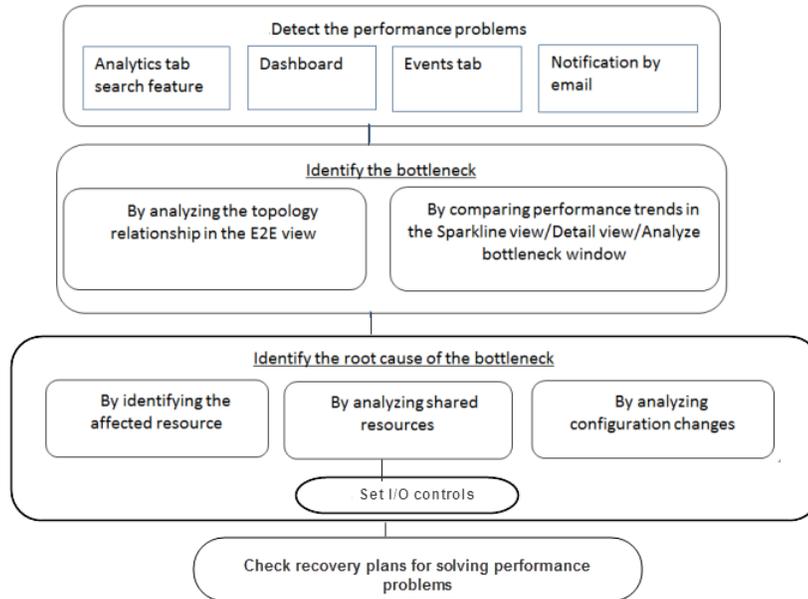
Component	Performance problem	Key performance metrics		
		Identify the resources with SLO violations or performance problems	Identify the related resources used by the affected resources	Identify the resources that might be the root cause
Server	CPU contention	VM <ul style="list-style-type: none"> ▪ vCPU Ready ▪ vCPU usage ESX <ul style="list-style-type: none"> ▪ pCPU usage ▪ Host CPU Ready¹ 	ESX <ul style="list-style-type: none"> ▪ pCPU usage ▪ Host CPU Ready¹ 	VM <ul style="list-style-type: none"> ▪ vCPU Ready ▪ vCPU usage
	Memory swap	VM <ul style="list-style-type: none"> ▪ Usage % ▪ Active memory ▪ Swap in/out rate ESX <ul style="list-style-type: none"> ▪ Swap in/out rate¹ 	ESX <ul style="list-style-type: none"> ▪ Usage % ▪ Active memory¹ ▪ Swap in/out rate¹ 	VM <ul style="list-style-type: none"> ▪ Usage % ▪ Active memory ▪ Swap in/out rate¹

Component	Performance problem	Key performance metrics		
		Identify the resources with SLO violations or performance problems	Identify the related resources used by the affected resources	Identify the resources that might be the root cause
	Memory contention	VM <ul style="list-style-type: none"> Balloon VM <ul style="list-style-type: none"> Balloon¹ 	ESX <ul style="list-style-type: none"> Usage % Balloon¹ 	VM <ul style="list-style-type: none"> Usage % Active memory Balloon
	Response time decrement	ESX <ul style="list-style-type: none"> pCPU usage Device Latency (R/W)¹ 		
Storage	Response time decrement	VM <ul style="list-style-type: none"> Latency (R/W) Hypervisor <ul style="list-style-type: none"> Latency (T)¹ LU (Volume) <ul style="list-style-type: none"> Response time (R/W/T) 	Port <ul style="list-style-type: none"> usage Processor <ul style="list-style-type: none"> MPB utilization¹ Cache <ul style="list-style-type: none"> Write Pending % Side file % Pool <ul style="list-style-type: none"> Utilization Parity Group <ul style="list-style-type: none"> Utilization % Read Hit % 	VM <ul style="list-style-type: none"> Read (KBps) Write (KBps) Read Operations Write Operations LU (Volume) <ul style="list-style-type: none"> IOPS (R/W/T)
Network	Error packet	VM <ul style="list-style-type: none"> droppedRx droppedTx Transmitted/received (KBps) 		VM <ul style="list-style-type: none"> Transmitted/received (KBps) PacketsTx¹ Packets Rx¹

¹ The performance metric is available in Analyzer detail view.

Analytics workflow

The workflow for analyzing performance problems and identifying the root cause is as follows:



1. [Detect performance problems \(on page 44\)](#)
2. [Identify the bottleneck \(on page 45\)](#)
3. [Identify the root cause of the bottleneck \(on page 50\)](#)
4. [Set I/O controls](#)
5. [Check recovery plans \(on page 53\)](#)

Detecting performance problems

View the threshold violations using the Dashboard tab and Events tab. You can configure the system to send email notifications when the threshold values are exceeded. You can also use the search feature in the Analytics tab to find the target resources for performance analysis.

Dashboard

The dashboard displays when you log on to Ops Center Analyzer. You can create a custom dashboard, and choose to view the reports of monitored resources.

The dashboard displays summary reports for the monitored resources, system and resource events, event trends, and consumer groups. The report widgets display the threshold violations and critical alerts detected on all monitored resources when threshold values are exceeded.

In the following dashboard view, the warnings display on the monitored VMs and volumes. From the report widgets, click links to access the E2E view and analyze the cause of the threshold violations.



Events tab

The Events tab displays a list of resource and system events. View the severity of each event, date and time of the occurrence, category, device, and the component name. You can navigate from the Events tab to the E2E view for further analysis.

Email notifications

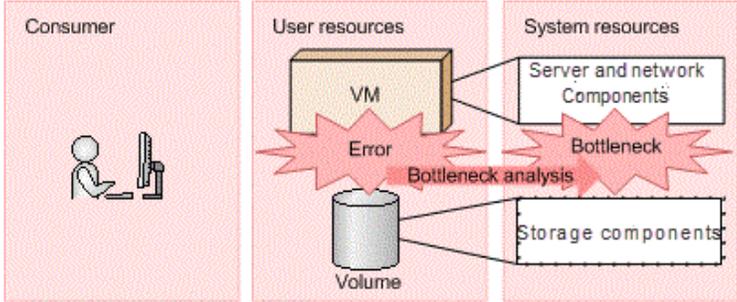
Ops Center Analyzer allows you to configure email notifications. When the threshold values are exceeded, the system sends an email to notify you of the potential performance problem.

Search

Use the search feature on the Dashboard tab to search for a resource in the Consumers, Servers, Switches, Storage Systems, and Volumes categories. From the returned search results, select the resources to analyze, and launch the E2E view or Sparkline view for further analysis.

Analyzing performance bottlenecks

The performance degradation in the user resources is caused by a performance bottleneck on the server, network, or storage components.



A performance bottleneck can occur for various reasons, such as CPU contention, inefficient load balancing, applications sharing storage pools, port and parity group utilization in shared infrastructure, cache utilization, changes in dynamic tiering policies, and configuration changes.

Identify and analyze the component causing the bottleneck in any of the following views:

- E2E view
- Analyze bottleneck window (except VSS Block)
- Sparkline view (except VSS Block)
- Detail view

Identifying the bottleneck in E2E view

Analyze the configuration of the infrastructure components in the E2E topology view.

The following procedure describes the workflow of tasks for troubleshooting the problem that occurred in a VM component.

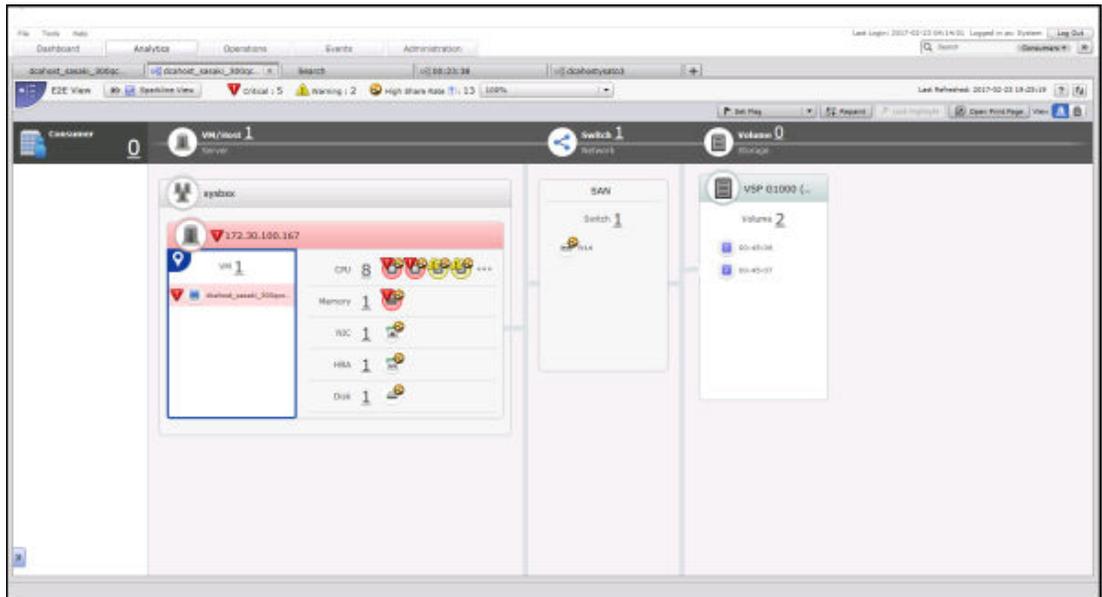
Procedure

1. Navigate to the E2E view in one of the following ways:
 - On the **Analytics** tab, perform a search for the target resources, and from the returned search results, select the target resource for analysis and click **Show E2E View**.
 - On the dashboard, the report widgets display the number of affected resources that exceeded the threshold values. For example, the VMs/Host report displays the number of affected VMs, and when you click the number link, a new window displays a list of monitored VMs. Select the resources to analyze and click **Show E2E View**.

2. In the **E2E View**, analyze the server-related and storage-related components to identify the resources causing performance problems.

- Click Server View  to get a server-oriented view of the business system configuration. The key components to monitor while analyzing the server performance are:
 - CPU
 - Memory
 - NIC
 - HBA
 - Disk
- Click Storage View  to get a storage-oriented view of the business system configuration. The key components to monitor while analyzing the storage performance are:
 - Port
 - Processor
 - Cache
 - Pool
 - Parity Group

In the following example, the alert indicators display on the VMs. When you analyze the VMs in the Server View, you can view alerts associated with the CPU server components.



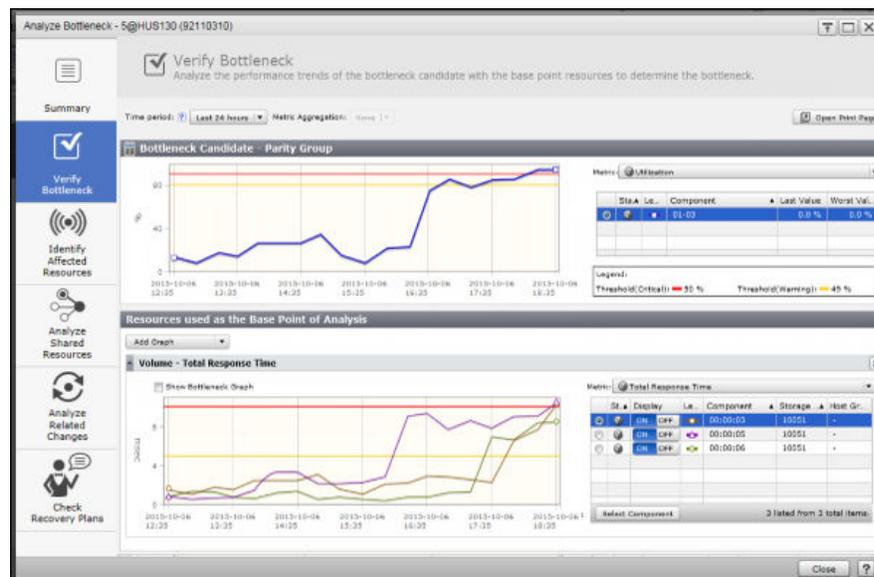
Note: The **E2E View** does not show component information such as CPU, memory, and disk space for Windows and Linux hosts.

- To analyze the resource associated with an alert, click the resource icon and select **Verify Bottleneck**.

For example, to check whether CPU is the bottleneck candidate, click the CPU icon and select **Verify Bottleneck**.

- The **Analyze Bottleneck** window displays the performance trend reports of the bottleneck candidate and the resource used as the base point of analysis. If the performance charts display similar trend patterns in the same time period, you can assume that the selected resource is the bottleneck candidate. If not, repeat the analysis for other resources with alerts in the **Verify Bottleneck** window.

For example, in the **Verify Bottleneck** window, the bottleneck candidate (CPU) appears in a graph in the upper pane, and the **VM** where the problem occurred appears in a graph in the lower pane. The performance charts display similar trend patterns in the same time period, which confirms that the CPU is the bottleneck candidate.



- You can also use the resource sharing percentage to identify the bottleneck in the shared infrastructure. On the E2E view tool bar, from the **Configuration Information** menu, select an option that displays the highest resource sharing percentage. Hover over the icons in the E2E view to see the shared percentage of each resource. Resources with a high sharing rate are the potential bottleneck candidates.

For example, when no alerts display on the dashboard, or in the E2E view, and when you cannot identify the bottleneck candidate using the **Analyze Bottleneck** window, use the resource sharing percentage to identify the affected resources in the shared infrastructure.

Select **Configuration Status** from the **Configuration Information** menu, to see information about the drives in the parity group, such as **Drive Status** and **Used Spare Count**. This tool is useful for parity groups with flash drives. Click on the **Parity Group** resource icon from the **E2E View** and select **Show Detail** to see details for the **SSD Used Endurance Indicator** and the **FMD Battery Life Indicator**.

Comparing performance trends in Sparkline view

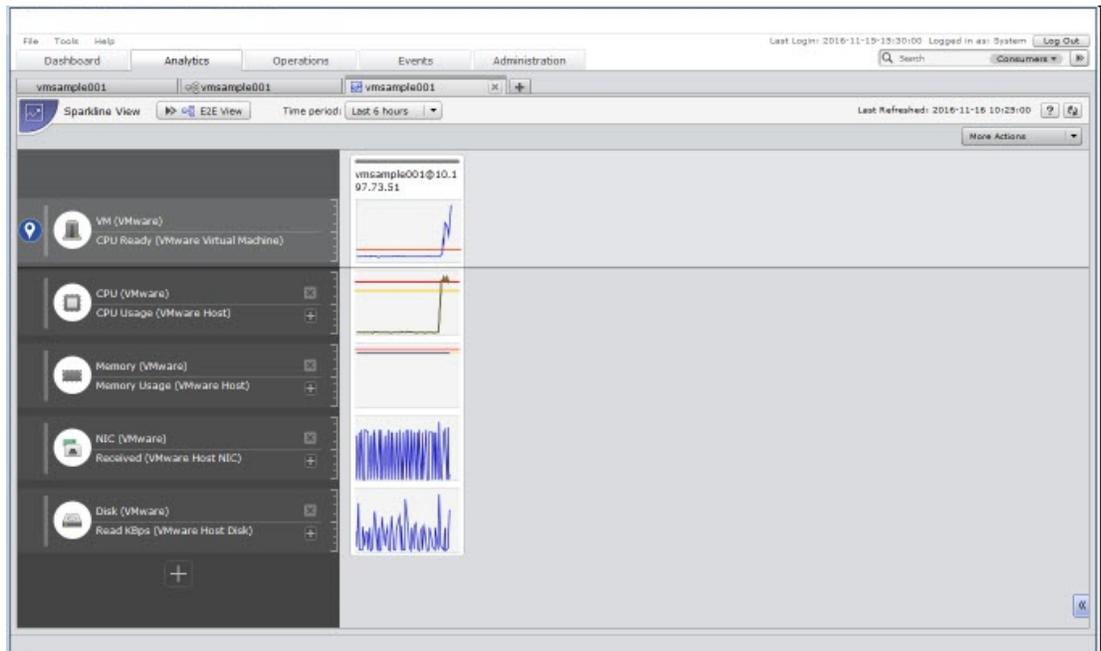
Use the Sparkline view to analyze the health and performance of the resources in your monitoring environment. The Sparkline view displays performance reports for multiple nodes in the same pane for a quick comparison between different nodes. Display detailed performance metrics for each node and find the correlation with other nodes.

The following procedure describes the workflow of tasks for troubleshooting the problem that occurred in a VM component.

Procedure

1. Navigate to the Sparkline view in one of the following ways:
 - On the **Analytics** tab, search for the target resources, and from the returned search results, select the target resources for analysis and click **Show Sparkline View**.
 - From the **E2E View**, select a resource, and then from the tool bar, click **Sparkline View**.
2. In the Sparkline view, analyze and compare the performance trends of the target and related resources. You can select more than one target resource for analysis.

In the following example, analyze the performance trends of the VM and the associated server components. The VM and CPU display similar trends in the same time period, which confirms that the CPU is the component affecting the performance of the VM.



**Note:**

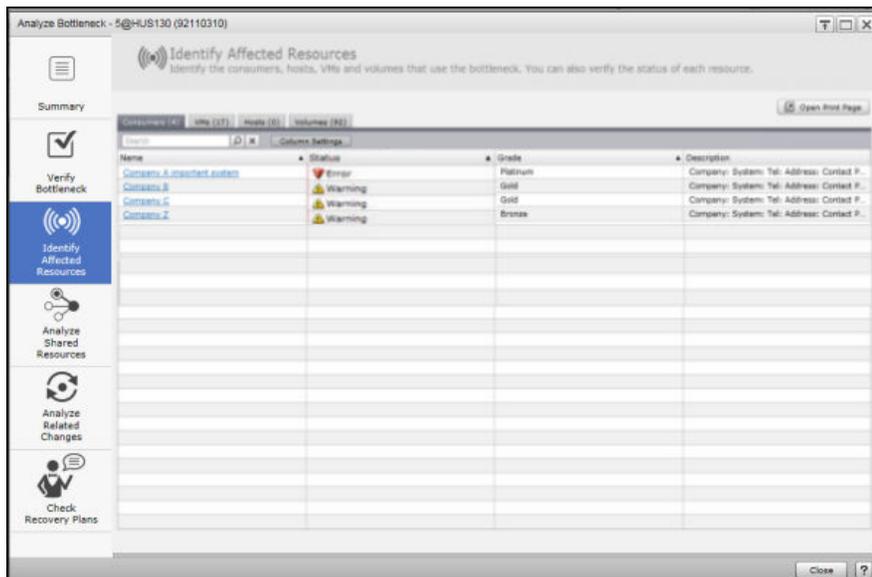
- The Sparkline view does not display the overall performance trends of Hypervisors and Storage Systems. However, you can analyze the performance of all associated components of hypervisors (such as, CPU, Memory, NIC, HBA, Disk) and storage systems (such as, port, processor, cache, pool, parity group) in the sparkline view.
- You can add a maximum of 20 target resources (base point resources) in the Sparkline view, and display a maximum of 200 graphs.

3. To analyze data with finer granularity, select the graph, and click **Show Performance** from the menu.

The performance window displays trends for the selected components.

Identifying affected resources

In the Analyze Bottleneck window, click the Identify affected resources tab. In this window, you can identify the consumers, hosts, VMs, and volumes that use the bottleneck candidate. You can also verify the status of each resource. Based on the severity level displayed, you can troubleshoot the performance problems associated with the resources.



Analyzing the cause of the bottleneck

The root bottleneck cause can be resource contention issues in the shared infrastructure, or configuration changes in the environment.

Analyzing shared resources

Performance problems arise when an application or a resource uses most of the available resources in the shared infrastructure. Ops Center Analyzer supports efficient optimization of the shared infrastructure by quickly identifying the resource contention problems.

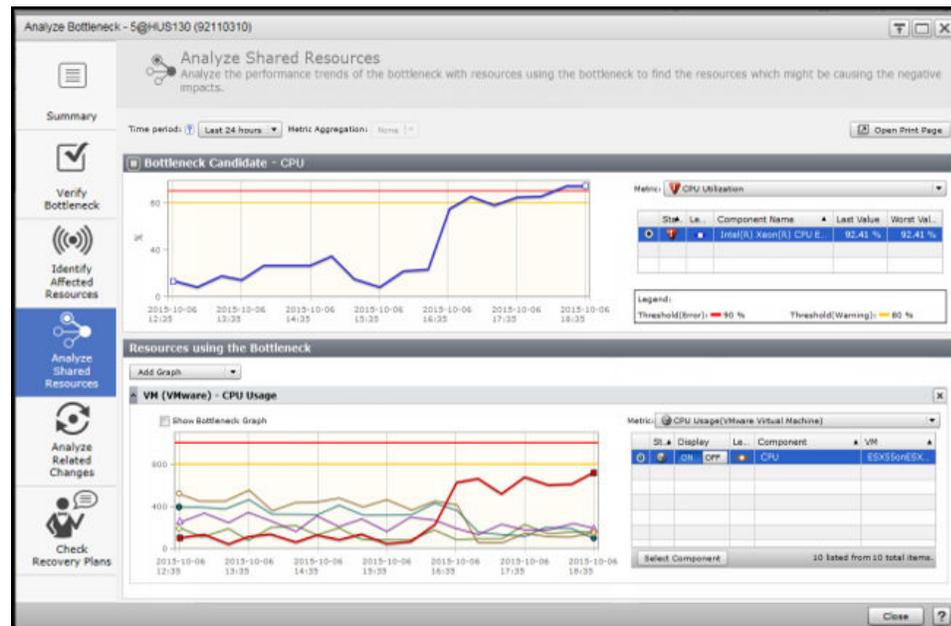
Procedure

1. In the **Analyze Bottleneck** window, click **Analyze Shared Resources**.
2. In the **Analyze Shared Resources** window, the performance charts appear for the bottleneck candidate and the resources using the bottleneck candidate.
3. Compare the performance trends of the bottleneck candidate with the resources that use it. If the compared resources display similar trends, then you can assume that the resource in the shared infrastructure is causing the bottleneck.

If you cannot determine the cause of the bottleneck from the displayed graphs, check the graphs of different metrics. To view the graph of a different metric, select it from the **Metric** list. You can also add a performance graph of a different component from the **Add Graph** menu.

In the following example, the CPU utilization value of a virtual machine is compared with the CPU utilization of other virtual machines and volumes. The performance trends confirm that one of the VMs in the shared infrastructure is overutilizing the CPU.

The CPU bottlenecks occur when several VMs run on the same physical machine, and end up sharing the same CPU. If the VMs (logical resources) share the same CPU (physical resource) and if one of the VMs utilizes the CPU more than the others, the total efficiency of the shared resource is degraded and the CPU utilization rate increases. The CPU might become saturated with requests because of resource contention problems.



Analyzing configuration changes

Ops Center Analyzer supports the tracking of infrastructure configuration changes. Analyze these changes and correlate them with the performance data to determine the effects of configuration changes on system performance and behavior.

Procedure

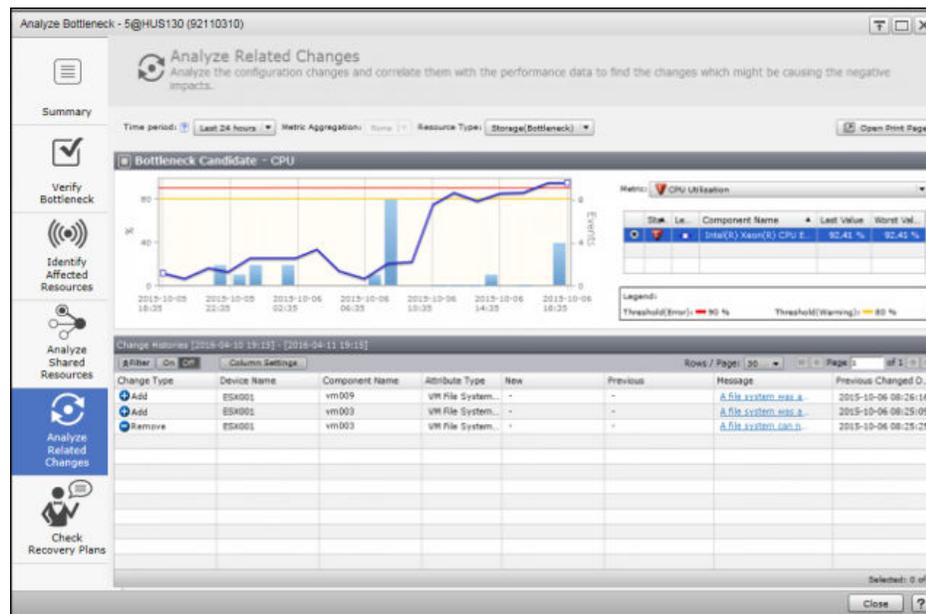
1. In the **Analyze Bottleneck** window, click **Analyze Related Changes**.
2. In the **Analyze Related Changes** window, analyze whether the bottleneck is caused by the changes in the system configuration.

The **Analyze Related Changes** window displays a combination chart that combines the features of the line and bar chart. The line indicates the performance of the bottleneck candidate and bars indicate the configuration change events. Analyze the change events to see if any of them caused performance variations in the bottleneck candidate.

The Change Events table in the lower pane displays a list of configuration-related changes and their details. Zoom in on the performance trend chart to select a shorter time period and view the change events that occurred in the selected time range.

In the following example, the performance data of the bottleneck candidate (CPU) is compared with the change events that occurred in the specified time period. You can correlate the performance data of the CPU and the change events to determine the effects on system performance. Based on the analysis, you can confirm that a lot of configuration change events caused performance degradation in the CPU.

To print the list of configuration change events, click the **Open Print Page** button. For a large number of events, the print operation might become unstable. From the **Rows / Page** pull-down menu, reduce the number of listed events, then retry the print operation.



Checking recovery plans

Ops Center Analyzer supports generating recovery plans for the processor and cache bottlenecks. The recovery plans provide guidance for solving performance problems. The recovery plan generation is supported for VSP 5000 series, VSP E series, VSP F series, VSP G series, VSP N series, VSP, and HUS VM series storage systems.

Procedure

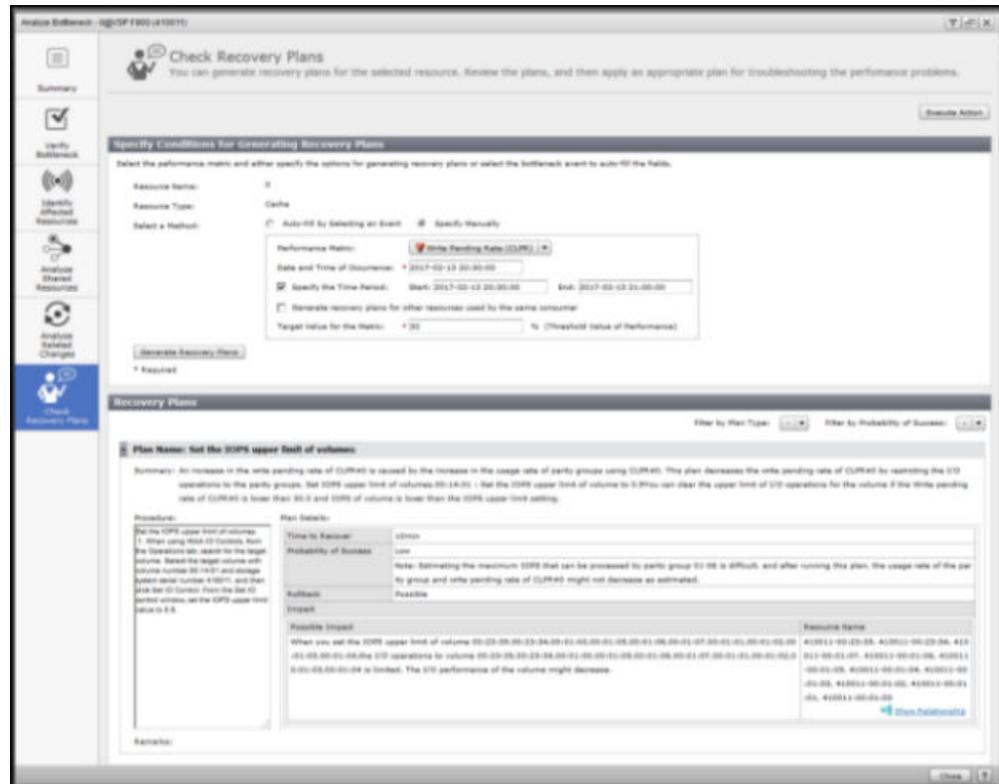
1. From the E2E View, select a resource you want to analyze and click **Analyze Bottleneck**.
 2. In the **Analyze Bottleneck** window, click the **Check Recovery Plans** tab.
 3. Select one of the following options to generate a recovery plan:
 - **Auto-Fill by Selecting an Event:**
 - Click **Browse Events** to select an event. The metric that exceeded the threshold value and the event date and time are auto-filled when you select a bottleneck event.
 - Specify the target value for the metric.
 - **Specify Manually:**

You must specify the following options manually:

 - Select the metric that exceeded the threshold.
 - Specify the date and time of the occurrence.
 - Specify the target value for the metric.
-  **Note:** Ops Center Analyzer supports generating recovery plans only for the following components and associated key performance metrics:

 - Processor: Utilization (MPB)
 - Cache: Write Pending Rate (CLPR)
4. (Optional) Select **Generate recovery plans for other resources used by the same consumer** to generate recovery plans for the bottleneck resource, and other resources used by the same consumer.
 5. Click **Generate Recovery Plans**.

View the system-generated recovery plans for the selected resource.



- Review the recovery plans and apply an appropriate plan to solve performance problems.

Executing actions

Execute predefined actions on a resource node. These predefined actions allow you to notify the appropriate IT administrators to troubleshoot the resource node problems.

Access the Execute Action window at any point during analysis from the following windows:

- E2E View
- Sparkline View
- Analyze Shared Resources
- Predictive analytics risk reports
- Check Recovery Plans
- Tools

Configure the system to perform the following actions:

- Send email notifications

For example, if a problem arises on a resource node, configure the system to send email notifications to the concerned administrator to troubleshoot the problem. Include the resource node information, troubleshooting methodology, and other information in the email template.

- Run a command

For example, if a problem arises on a resource node, execute the command action to automatically register the problem in a trouble-ticketing system, so that the concerned IT administrator can troubleshoot the node.

- Submit a service through Ops Center Automator.

For example, submit a service based on a template from Ops Center Automator to resolve a recurring problem.

- Start granular data collection.

For example, if your analysis requires data collection in intervals less than a minute, run the command for granular data collection.

Procedure

1. From one of the entry points, click a resource icon, and then select **Execute Action**.
2. In the **Execute Action** window, view the list of defined actions.

The Actions list will be empty if you have not defined any actions previously. Following is the high-level procedure for creating event actions:

- a. Create an event action definition file and save it to the Ops Center Analyzer installation folder.

Ops-Center-Analyzer-installation-folder/Analytics/conf

The default installation folder is `/opt/hitachi`.

- b. Restart Ops Center Analyzer, or run the `reloadtemplate` command to view the defined actions.

3. Under Actions, select an action.

- To send email notifications, select the action type Mail and then click **Launch Editor**.

The email template is launched in the email editor. Alter the email template to suit your requirements.

- To run a command, select the action type Command.

The details of the command are displayed. Edit the **Command Arguments** field, then click **Execute Command**. If you want to run more than one command at the same time, you can create a batch or shell command. In the **Execute Command** window, review the details of the command and then click **OK**.

- To collect data at shorter time intervals, select **Granular Data Collection** or **Granular Data Collection with Volume**.

Submitting services through Ops Center Automator

Submit a service in Ops Center Analyzer based on built-in action templates or Ops Center Automator service templates.

In addition to running actions (such as email and CLI commands) through the Execute Action feature, you can submit services through Ops Center Automator when it is linked to Ops Center Analyzer. Some action templates already exist in Ops Center Analyzer. You can also access service templates available in Ops Center Automator.

Before you begin

- Ops Center Automator must be linked to Ops Center Analyzer through the Common Component.
- User permissions must be set for both Ops Center Analyzer and Ops Center Automator user accounts.
- The definitions in the Ops Center Analyzer action templates must match those in the corresponding Ops Center Automator service names and service group names.

Procedure

1. Click on the target resource and select **Execute Action**.
2. In the **Execute Action** window, select the action or service template you want to submit, then do the following:
 - Scroll through the list of default action templates.
 - Check the box to see available service templates in Ops Center Analyzer
3. Click **Launch service execution**, review the parameters in the **Settings** field, then click **Submit**:
 - If you are running a preexisting action, the parameters are autopopulated.
 - If you are running a service, you might need to enter the parameters manually.
4. To monitor the task, click **Launch Ops Center Automator** from the **Execute Action** window.

Granular data collection

Use the granular data collection feature to analyze second-level data.

Ops Center Analyzer provides numerous reports on system and user resources. Typically, system administrators manage resources for application data, and dashboard reports give sufficient information on system resources (such as volumes) that support a widely used application. When a bottleneck occurs, system administrators look for data spikes. However, this data might not be available on the dashboard because of the time delay between monitoring and reporting. While this report data displayed on the dashboard is collected and delivered within minutes of the events, in this situation, system administrators require second-level data. This data is collected in intervals of seconds, providing the granularity necessary to search for data spikes.

When analyzing storage resource bottlenecks, the root cause might be one of the following:

- Port saturation
- MPB utilization
- High response times for volumes

Depending on the way granular data collection is run, the CSV files from the resulting output consist of performance statistics for all ports, volumes, processor data (MP), and activity-rate ranking for each processor allocated to an MP blade and volumes of that storage system.

Persistent performance issue

A critical event occurs multiple times throughout the day, but you do not know what component in storage resources is affecting the user resource, such as an application. The following workflow shows the stages of analysis you follow to determine the root cause:

- Start searching through candidate volumes or the volume listed in an SLO warning.
- Analyze the E2E View to identify the resources associated with the Consumer.
- Analyze the Sparkline view to identify the spike in the trend patterns.
- Identify affected resources to determine the bottleneck resource.
- Run granular data collection for that resource, for example, from 15 minutes to an hour.
- Analyze CSV output.

Typically, this use case requires that you run granular data collection for a longer duration since the spikes might not occur predictably.

Performance issue that occurs consistently during a specific time

If you notice a performance problem in storage resources consistently at the same time, for example, from 1 to 3 PM during the week, you can run granular data collection on various components for as long as duration of the period in which the problem occurs (in this example, two hours).

Typically, you run granular data collection many times for shorter periods in this use case.

Collecting granular data

The Granular Data Collection tool enables you to collect data in intervals as short as one second for analysis.

In cases where detecting spikes in trend patterns is necessary, data collection in intervals of minutes might not provide sufficient granularity of data. To collect data in intervals of seconds, invoke the Granular Data Collection tool, which generates output to CSV files. You can upload this data to a spreadsheet or charting application for graphical analysis of data spikes in the trend patterns.

Before you begin

- Verify the following OS support:
 - The Ops Center Analyzer server is running Linux OS.
 - The RAID Agent (or Tuning Manager - Agent for RAID) host server is running a version of Linux OS supported by Ops Center Analyzer.
- The target storage system is a model in the VSP 5000 series, VSP E series, VSP F series, VSP G series, VSP N series, VSP, and HUS VM families.
- The version of Tuning Manager - Agent for RAID is 8.5.1 or later.

Procedure

1. Select a storage resource or storage component, and then select the **Execute Action** option from the menu.

During analysis, select any of the following resources and components:

- storage system
- volume
- port
- processor
- pool
- parity group
- cache



Note: Other action templates similar to the Granular Data Collection tool might display in the **Execute Action** window.

2. On the **Execute Action** window, select **Granular Data Collection**. To run the tool on a specified volume, select **Granular Data Collection with Volume**. After selecting either tool, specify the following arguments and click **Execute Command**.

All arguments are optional except where indicated that they are mandatory.

- a. Interval of data collection, in seconds (1-60):
 - `-intervalInSec`
 - Default value: 1
- b. Period of data collection, in seconds (1-3600), for one file:
 - `-periodInSec`
 - Default value: 300
- c. Number of times (1-24) data collection runs:
 - `-repeatTimes`
 - Default value: 1
- d. Starting time of data collection using any time expression (supported by date command usage in Linux OS), such as, 13:00, tomorrow 13:00, 1 PM, or now:
 - `-startTime`
 - Default value: now

- e. Timezone offset (-2359 to -0000 and +0000 to +2359) between the user client and UTC:

`-timezone`

Default value: +0000

- f. (Mandatory) The DKC model name of the target storage system:

`-model`

Default value: Specified automatically by the selected resource.

- g. (Mandatory) Serial number of the target storage system:

`-serialNumber`

Default value: Specified automatically by the selected resource.

- h. Number of the volume that you selected at the entry point. This parameter only displays when the action is **Granular Data Collection with Volume**:

`-ldev`

Default value: Specified automatically by the selected resource.

You can run granular data collection on multiple volumes (128 maximum) by concatenating volume numbers with commas.

- i. The folder name that is the destination folder of the output files:

`-pathLabel`

Default value: `result`

Example of arguments in one string:

```
-intervalInSec 1 -periodInSec 300 -repeatTimes 1 -startTime "now"
-timezone "+0000" -model "VSP G1000" -serialNumber "345678" -ldev
"01:23:45" -pathLabel "result"
```

Guidelines for granular data collection:

- For the first second, the acquisition period is less than 1 second, and if there is no I/O in that period, the performance value is 0.
 - Granular data collection cannot be cancelled once you run the command.
 - You can only run the command for granular data collection one instance at a time.
 - Running granular data collection might affect the performance of storage resources if overused. When running granular data collection, monitor the following:
 - The duration in which the MP blades and MP units are used exclusively depends on the number of volumes for which the data is obtained and on the collection interval. Therefore, keep the number of volumes to a minimum.
 - Because a small collection interval increases the load on the MP blade and MP unit, specify an appropriate value for the collection interval option. If a load problem on the MP blade or MP unit persists, consider allocating the command device that the RAID Agent (or Tuning Manager - Agent for RAID) accesses to other MP blades or MP units.
3. View the results of the granular data collection task in CSV files located in the following directory: `/Ops-Center-Analyzer-installation directory/Analytics/webapps/webroot/Analytics/result/GranularData/start_date/label/storage_id/secdata/yymmdd/`.

Where:

- *start_date with timezone offset* is the date that the data is collected with the time zone offset in the `yyyymmdd[+/-]hhmm` format.
- *label* is the label defined by the user in the `pathLabel` option.
- *storage_id* is the target storage system identifier: *model-name_serial-number*.
- *yyyymmdd* is the date that the granular data collection is run.

Output path example:

```
/Ops-Center-Analyzer-installation directory/webapps/webroot/  
Analytics/result/GranularData/20170119+0000/result/  
VSP_G1000_345678/secdata/20170119/
```

You can also view the results on a web browser:

```
http://host-name or IP-address of Ops-Center-Analyzer:port of  
Ops-Center-Analyzer server/Analytics/result/GranularData/  
index.html
```

Output web address example:

```
http://172.0.0.1:22015/Analytics/result/GranularData/  
20170119+0000/result/VSP_G1000_345678/secdata/20170119/index.html
```

4. CSV output is not deleted automatically. You can delete it manually. You can also set up routine file deletion, as shown in the following example.

Delete output after 3 days at 1:00 PM:

```
$crontab -e
```

```
01*** product installation folder/Analytics/sample/rmOldExportData.sh 3
```

Using Hitachi Ops Center Analyzer for data analysis: from deep dive to recovery planning

Hitachi Ops Center Analyzer provides an intuitive UI for performance monitoring, management, and troubleshooting. The Analyzer detail view collects data from monitored targets (such as storage systems, hosts, and switches) using software probes that support each device or environment. Analyzer detail view also provides historical trend analysis and extensive report generation capabilities.

This analytics solution provides end-to-end monitoring and troubleshooting capabilities for your infrastructure resources, from host to storage system. The basic workflow for Performance Analytics troubleshooting is called the MAPE loop:

- Monitor
- Analyze
- Plan
- Execute

When reviewing and evaluating reports and event information on the Ops Center Analyzer Dashboard, you can also perform a deep dive analysis by launching the Analyzer detail view UI. The deep dive is part of the **Analyze** segment of the MAPE loop workflow.

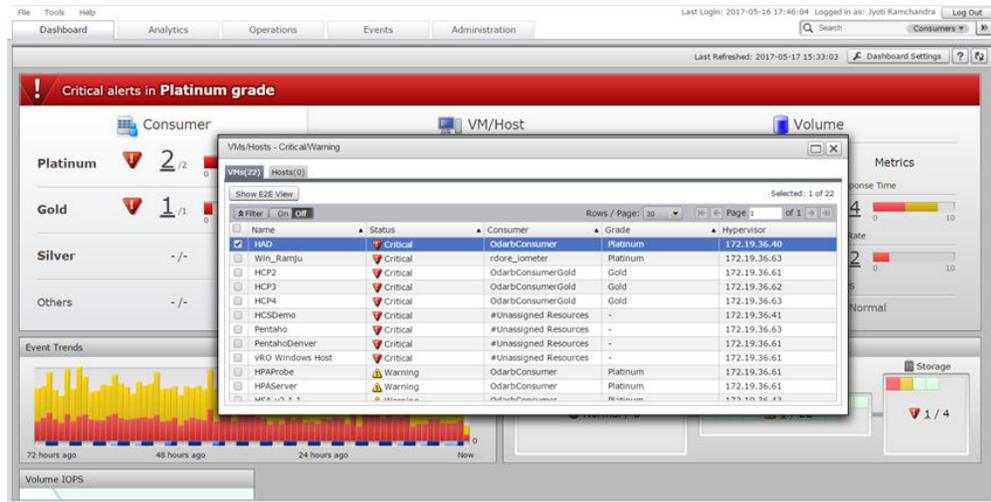
The following workflow is an example of how to use this troubleshooting methodology as an infrastructure administrator who manages user resources (such as consumers, VMs, and volumes) and system resources (such as cache, ports, CPUs, and disks).

Viewing the dashboard

As an infrastructure administrator, you set up dynamic thresholds on the user resources you are monitoring. After seeing nine critical alerts on VM/Host resource gauge, you become interested in troubleshooting a threshold violation.



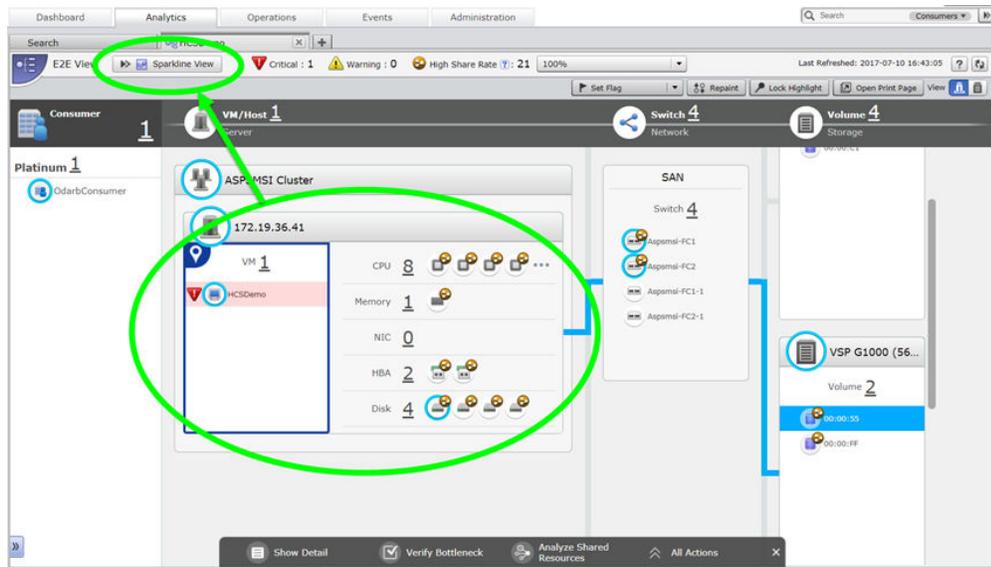
You browse the resources with critical alerts and select the target VM to analyze in the E2E View.



Using E2E or Sparkline views

The E2E view represents the topology of infrastructure resources: from host, to fabric switch, to storage system. The infrastructure administrator sets the base point of analysis on the target resource for analysis. This view enables you to see the relationship between resources.

To move deeper into the underlying resources, you can launch the Sparkline view, which presents multiple charts that track performance by component. Use this view to correlate performance trends between user and system resources.



Using additional troubleshooting tools

Ops Center Analyzer offers multiple troubleshooting tools for isolating a bottleneck candidate and identifying the root cause. You can launch any of the following tools for further analysis:

- **Verify Bottleneck:** Use at the initial stage of analysis to compare performance charts of the base point of analysis with the bottlenecked candidate.
- **Identify Affected Resources:** Use to display the user resources that rely on the bottlenecked resource.
- **Analyze Shared Resources:** Use if you suspect that the root cause of the problem is resource contention, a noisy neighbor that disrupts the balance of resource usage. You compare performance charts of the bottleneck candidate to the resources using the bottleneck. After comparing performance across a number of resources with Analyze Shared Resources, you isolate the actual bottleneck.
- **Analyze Related Changes:** Use if **Analyze Shared Resources** does not reveal the actual bottleneck (noisy neighbor), or if you suspect that the root cause of the problem is a recent configuration change. In this view, you compare performance charts with configuration events. The bar graph portion of the chart represents the configuration changes made at a particular time. You can click on a bar to list those changes.

Performing a deep dive analysis

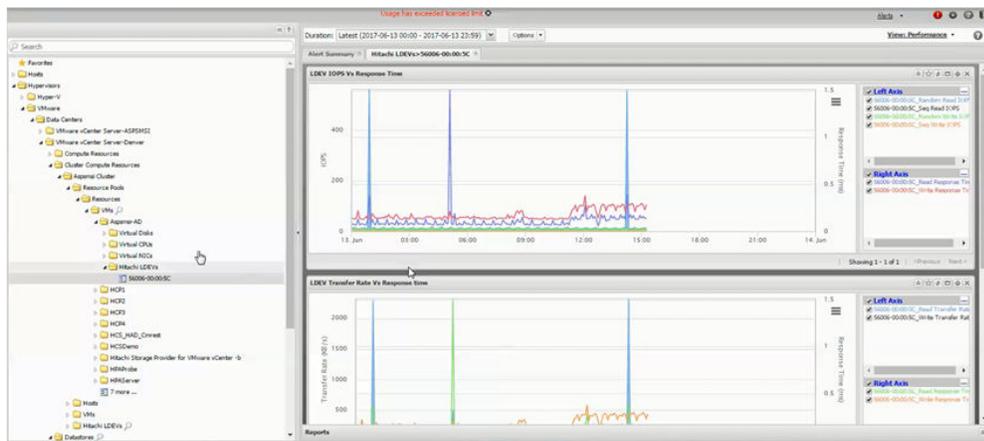
Regardless of which tool you use, after you have isolated the bottleneck candidate and validated the root cause, you can collect more information to understand its origin. For example, you have identified a storage system as the bottleneck. Subsequently, you want to understand how the problem affects other resources or vice versa. This phase of the troubleshooting analysis is called the deep dive. In a deep dive analysis, you can compare the data of various components from the resource tree, which displays all the resources and their components in your infrastructure, and run a customized report against that data.

To proceed with the deep dive for information, launch the Analyzer detail view UI, which provides detailed reports at the component level. You can launch this component-level view from the following windows in the Ops Center Analyzer UI during analysis:

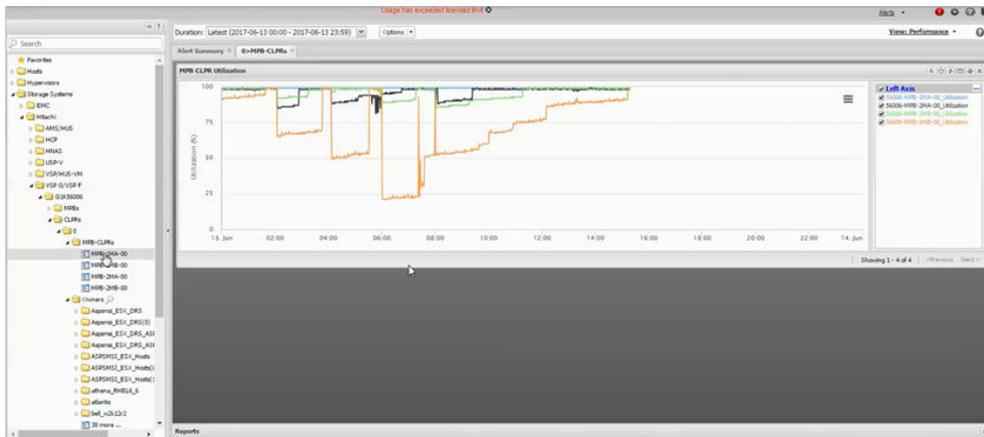
- E2E view
- Sparkline view
- Performance tab of the Show detail window for a resource
- Analyze Shared Resources
- Analyze Related Changes

When analyzing system resources in Analyzer detail view, you can view performance charts based on various metrics to correlate components with resource performance. For example, you have validated the root cause of the storage system bottleneck, but you want to perform further analysis in Analyzer detail view.

The following figure examines the performance of the volume from the VM side. This report, **LDEV IOPS versus Response Time**, displays spikes at specific times, which you can then use as reference points for when the I/O activity was particularly intensive during otherwise typical workloads.



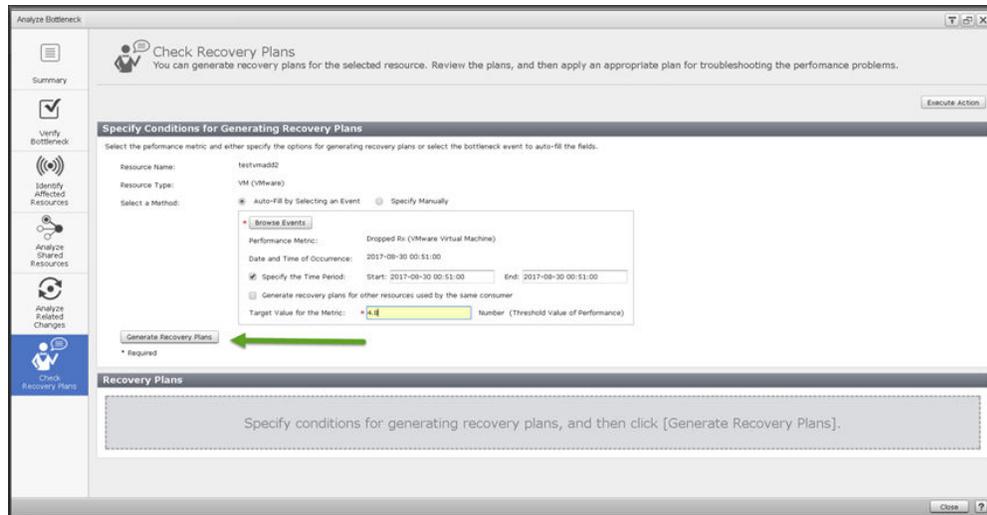
Digging deeper, you discover the storage systems and volumes associated with a particular VM. You cross-reference the resources in the VM performance chart and determine the component with the performance that correlates to the VM. In this example the resource that correlates with VM performance is the cache on the storage side (CLPR). This workload is typically intensive, but you realize that the times when the resource reached 100% correlate with the spikes in the **LDEV IOPS versus Response Time** report.



Often, the performance problem is a recurring trend; for example, when monitoring certain infrastructure resources, you notice spikes in I/O activity every weekday at 3 PM. When you create a customized report, you discover this trend has persisted for six months. (In theory, you can review performance from months to years.) This capability to review past performance adds a historical element to deep dive analysis.

Initiate recovery plan to solve the performance problem

After establishing the correlation between the two charts, you return to the Ops Center Analyzer UI to initiate a recovery plan. You can enter the key metric, date, and time of the problem occurrence, and the target value for the metric. In this case, the problem component is the CLPR; the key metric is IOPS. You can specify conditions, then review the recovery plan generated by Ops Center Analyzer before running it.



After the recovery plan runs successfully, you can adjust your thresholds with new metric settings to monitor the user resources (in this case, the VM and the affected volume). At this stage, you have completed the MAPE loop.

E2E infrastructure topology view

The E2E topology view provides the detailed configuration of the infrastructure resources and lets you view the relationship between the infrastructure components. You can manually analyze the dependencies between the components in your environment and identify the resource causing performance problems. By using the topology maps, you can easily monitor and manage your resources. Use this view to monitor resources in your data center, including applications, virtual machines, servers, networks, and storage systems.

In the E2E view, each node represents a resource, and the connecting links represent the relationship between the infrastructure components. You can analyze a target resource and all associated resources. You can also view alerts associated with all related resources and trace the problem at the root level. The node-based E2E view helps you analyze the problem on the affected node and its impact on other resources. You can also open the Analyzer detail view UI to view a detailed performance report for a selected resource.

Topology view components

The E2E view displays the topology related to the selected resources under the following default infrastructure groups:

- **Consumer:** The name of the consumer group to which the selected resource belongs and the details about the consumer grade level.
- **Server:** The associated server components, such as VMs and hosts.
- **Network:** The associated network components, such as switches.
- **Storage:** The associated storage components, such as volumes.

A number link is shown next to each resource icon. For example, when you select a storage subsystem as a target resource for analysis, and if 50 volumes belong to this storage subsystem, the value Volumes 50 is shown under the Storage infrastructure group. Click the Volumes link to open the Volumes - Storage window, which displays details about the volumes in the storage subsystem. From the Volumes list, select the priority of volumes that you want to analyze in the E2E view.

E2E view tool bar

The tool bar provides quick access to frequently used menu options and icons:

Options	Description
Sparkline View	Navigate to the Sparkline view to analyze the performance of the base point resource and the related resources to identify the bottleneck.
Critical	Number of critical alerts in the topology view.
Warning	Number of warnings in the topology view.
Configuration Information	Number of indicators for configuration information.
Configuration Status	Information about drives, such as availability or the battery life of an SSD drive.
Copy Pair Information	Copy pair information for volumes.
VSM Information	Virtual storage machine information for copy pair volumes.
High Share Rate	Resource sharing percentage of a shared resource. Hover over the resource icons to display the share percentage for each resource. Resources with high share rate are potential bottleneck candidates. The Share Rate value is not displayed when you set a Hypervisor or Storage System as the base point of analysis. Select OFF to turn off this feature.
Configuration Status	Information about drives, such as availability or the battery life of an SSD drive.

Options	Description
Storage and Server Views	<p>The following topology views are supported:</p> <ul style="list-style-type: none"> ▪ Storage View: Maps the storage-related components, such as ports, processors, cache, pools, and parity groups. ▪ Server View: Maps the server-related components, such as CPU, memory, NIC, HBA, and disk.
Lock Highlight	<p>Select a resource node and click Lock Highlight to highlight all related components in the topology view. The resource configuration remains highlighted until you release the lock on the resource node. This feature helps you understand the links between components and analyze the system configuration in detail.</p> <p>To release the lock, click Lock Highlight again.</p>
Repaint	<p>Select a resource node and click Repaint to move the resource from bottom-to-top or right-to-left to the prime position. Use this option to change the display order of resources.</p>

E2E view menu bar

Menu bar items	Menu items and description
Show Detail button	<p>Select a resource and click Show Detail. The performance summary report of the resource opens in a new window. You can also view the events related to the resource in the Events tab.</p>
Show Report in Analyzer detail view	<p>Click a resource icon and select Show Report in Analyzer detail view. The Analyzer detail view UI opens in a separate browser window. The resource tree opens to the selected resource, along with the latest available report in the Performance view.</p>

Menu bar items	Menu items and description
Analyze Bottleneck menu	<p>Select a resource and click Analyze Bottleneck. The Analyze Bottleneck Summary window opens. From the summary window, you can display the following tabs for the detailed analysis:</p> <ul style="list-style-type: none"> ▪ Verify Bottleneck ▪ Identify Affected Resources ▪ Analyze Shared Resources ▪ Analyze Related Changes ▪ Check Recovery Plans
Action menu	<ul style="list-style-type: none"> ▪ Change Base Point: Select a resource and click Change Base Point to change the node of analysis. The topology view opens for the selected resource in a new window. ▪ Execute Actions: Select a resource and click Execute Actions to run predefined actions on a resource node. The event actions allow you to send notifications to your administrator for troubleshooting performance problems.
Set Flag menu	<p>Select a resource and click Set Flag to flag a resource so you can analyze the flagged resource at a later point. To remove the flag, click Unset Flag.</p>
Show Prediction	<p>Select a resource and click Show Prediction to generate a report showing the predicted performance trend for that resource. The report is based on the risk profiles that you select. After the report is generated, go to the Predictive Analytics tab to view the results.</p>

Chapter 5: Analyzing performance trends with predictive analytics risk reporting

Use predictive analytics risk reporting tools to predict performance and utilization of infrastructure resources.



Note: Predictive analytics does not support Virtual Storage Software Block.

Predictive analytics risk reporting overview

With predictive analytics risk reporting, Ops Center Analyzer allows you to generate risk analysis reports based on infrastructure performance trends.

Predictive analytics risk reporting uses several predictive models to calculate the predicted performance trend. Once the predictive analytics license is enabled, the analytics engine uses historical data collected by the Analyzer probe server to calculate performance trends for your risk reports.

Anticipate performance trends

Anticipate performance trends across your infrastructure by generating risk reports from system-defined report definitions, which are available for immediate use.

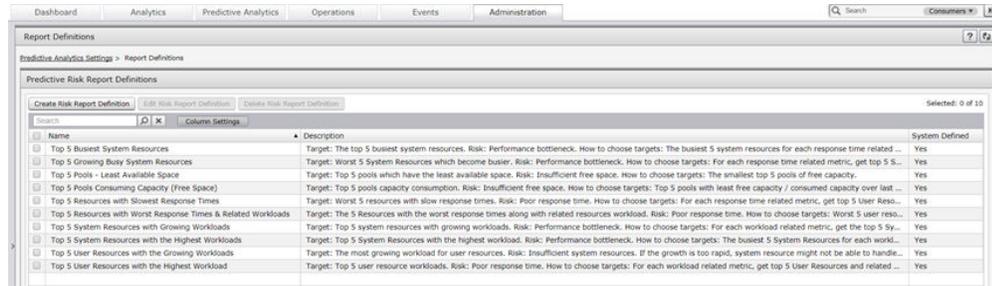
Use risk reports to report on the affected resources in your system. These resources are at risk of threshold violations. Ops Center Analyzer automatically tracks risk in affected resources and consumers by periodically calculating performance data trends.



Note: The calculation for selecting affected resources and consumers is executed when the Ops Center Analyzer starts. You can also set the time for these reports to run. The default starting time is 23:23 hours. It might take some time to complete the calculation after a fresh installation.

Generate risk reports on the following:

- Capacity shortages
- Capacity planning (growth)
- Slow response times
- Anticipated performance bottlenecks
- Risk for platinum consumers
- Growing workloads for user and system resources



About predictive risk profiles

Risk profiles specify the target metrics and time span for the predicted performance trend in your risk reports.

Ops Center Analyzer features default risk profiles for immediate use. These system-defined profiles cover a variety of use cases. On **Profiles Settings**, click on a profile and view the details provided on the **Profile Information** and **Target Metrics** tabs.



Note: Users cannot edit or delete system-defined risk profiles.

The **Create Risk Profile** function allows you to choose the parameters manually. You can also use the **Copy From** option to copy an existing profile from the menu, and then modify the settings according to your risk assessment goals. You can edit and delete user-created risk profiles from **Profile Settings**.

About predictive risk report definitions

Risk report definitions specify the target resources and risk profiles for generating a risk report.

You can select hypervisors, VMs, hosts, FC switches, storage systems, and volumes as target resources for the report. Use **Report Definitions** to choose a report definition, and view the details on the **Report Information**, **Profile Information**, and **Target Resources** tabs.



Note: Users cannot edit or delete system-defined report definitions.

You can also create your own risk report definitions. Use the **Create Risk Report Definition** function to choose the parameters manually. You can also use the **Copy From** option to copy

an existing report definition, and then modify the settings according to your risk assessment goals. You can edit and delete user-created risk report definitions from the **Report Definition**.

Predictive analytics risk reporting workflow

Calculate performance trend projections by choosing a report definition and adding risk profiles.

In the report definition, select resources for the predictive modeling calculations. Use the risk profile to define the target metrics and time period. The end result is a trend chart with performance trend projections for the selected resources. The reports are similar to the performance-by-metric reports in the **Analytics** tab. Ops Center Analyzer provides system-defined risk profiles and report definitions for immediate use. Follow the steps in this workflow to begin:

1. Verify licenses. Check the Ops Center Analyzer **Licenses** window to verify that the licenses are valid.
2. Set up risk profiles. These profiles establish the time period and the target metrics for the performance trend projection.
3. Select a risk definition. This definition instantiates the output for a performance trend projection by combining multiple risk profiles for target resources.
4. Generate a report. On the **Predictive Analytics** tab, choose a report definition. Click **Predict Result** to calculate the trend projection for that report definition.
5. Analyze the report. On the **Report Results** tab, choose a report. The selected report opens in a new tab.
6. Avert the risk. If the performance trend projection shows a potential threshold violation or degraded performance, access the **More Actions** menu to mark a resource for follow-up analysis, send an email, or run a script. Likewise, you can take corrective action by using available tools in Ops Center Analyzer, such as assigning resources, setting I/O controls, or generating recovery plans.

Adding the Predictive Analytics license

To use risk reporting in Ops Center Analyzer, you must have a valid Predictive Analytics license for each storage system in your environment.

Obtain a Predictive Analytics license from your Hitachi Vantara representative.

Procedure

1. In the **logon** window, click **Licenses**.
2. Use either of the following methods to register the license:
 - Enter the license key.
 - Specify the license file.
3. Click **Save**.
The license is added to the list.

Creating a predictive risk profile

Create a risk profile for projecting performance trends from 24 hours to 6 months in the future.

Before you begin

You must have a valid Predictive Analytics license to use the predictive analytics risk reporting feature.

Procedure

1. Click the **Administration** tab, then select **Predictive Analytics Settings > Profile Settings**.
2. Click **Create Risk Profile**.
3. On the **Copy From** menu, do one of the following:
 - Select a system-defined or an existing user-defined profile. You can rename it, then modify the parameters according to your risk assessment goals.
 - Select a blank profile (-), the default option in the menu, and construct a new profile by manually selecting the parameters.
4. Name the profile (required) and provide a description (optional).
5. Select a time period.
For example, select **24 hours** or **7 days** for near-term performance trends. For a long-term projection to determine utilization trends, select **1 month** or **6 Months**.
6. In the **Target Metrics** table, add or remove performance metrics to or from the profile. Click **Add Metric**, make the following choices, and click **OK**:
 - a. Use the **Resource Type** menu to display a list of metrics by resource type.
 - b. Select target metrics in the **Available Metrics** field, then click **Add**. When the **Selected Metrics** field is complete, click **OK**.
7. When the Target Metric table is populated, add a monitoring threshold for each target metric by adjusting the **Monitoring setting (ON/OFF)** and entering a threshold value. Some threshold values have predefined settings, which are adjustable.
8. Click **Save**.

Creating risk report definitions

Create risk report definitions to view performance projections for selected target resources.

Before you begin

You must have a valid license to use the Predictive Analytics risk reporting feature.

Procedure

1. Click the **Administration** tab, then select **Predictive Analytics > Report Settings**.

2. On the **Copy From** menu, do one of the following:
 - Select a system-defined or an existing user-created report definition. Rename it, then modify the parameters according to the risk assessment goals.
 - Select a blank profile (-), the default option, and construct a new profile by manually selecting the parameters.
3. Name the report definition (required) and provide a description (optional).
4. Select one or more risk profiles.
5. Specify the resource scope:
 - To include only selected target resources, choose **Selected**.
 - To include all resources associated with the selected target resources, choose **All Associated**.
6. Click **Add Resources**.
 - a. On the **Add Resources** window, select **Resource Type** from the menu.
 - b. Select **Available Resources** and add **Selected Resources**. You can also remove resources from the list.
 - c. Click **OK**.
7. Go to the **Predictive Analytics** tab to generate the report.

Generating a risk report

Specify risk profiles in a report definition to generate performance trend projections.

Before you begin

- You must have a valid Predictive Analytics license for the predictive analytics risk reporting feature.

Procedure

1. Generate a report in any of the following ways:
 - From the **Predictive Analytics** tab, click the **Report Definition** subtab, select one or more report definitions, then click **Predict Result**.
 - From the **Predictive Analytics** tab, click the **Search** subtab. Select a target resource category, such as consumers, servers, switches, storage systems, or volumes, then enter a resource name. Click **Show Predictive Report**.
 - From the **E2E View** window, select one or more consumers and resources, then click **Show Predictive Report**. The **Select Profiles** dialog opens.
 - a. Specify the **Target Resource** scope:
 - **Selected resources** includes only the resources you select.
 - **All resources associated with the selected resources** includes all associated resources in addition to resources you select.
 - b. Select one or more profiles from the **Predictive Risk Profile** menu, then click **OK**.

The generated report opens in a new subtab on the **Predictive Analytics** tab. View report execution status on the **Report Results** tab.

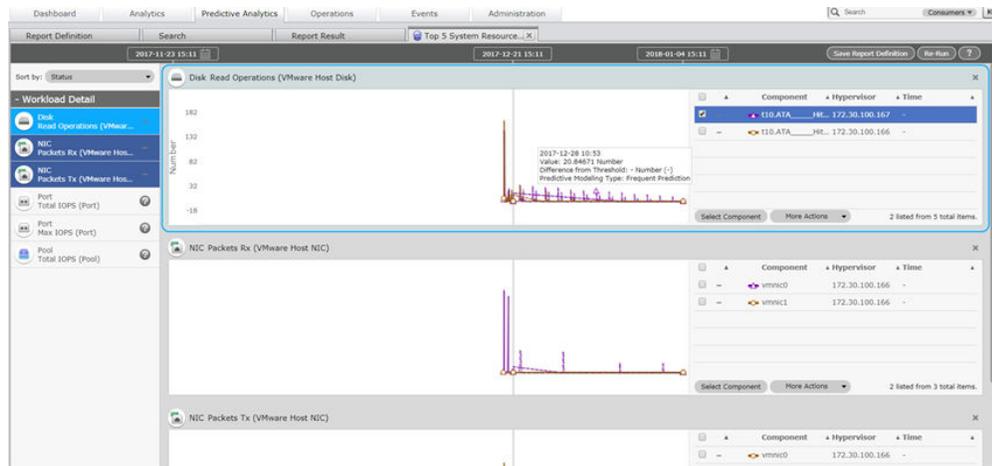
2. On the **Report Results** subtab, select the generated report to be opened, and click **Show Result**.
If you generated a report from the **E2E View** window or **Search** subtab, the report name begins with *OnDemand* followed by the timestamp.
3. On the left pane of the report, click a metric to view the predicted performance trend.

Understanding performance trend projections in risk reports

You can gain insight into near-term trends and long-term usage of your infrastructure resources from predictive analytics reports.

Report contents

The report is similar to other performance reports in Ops Center Analyzer and includes twin graphs. The left side represents current performance, and the right side represents the predicted performance trend. Zoom in on the graph by clicking and then dragging the cursor over an area. Double-click the graph to return to the normal view.



Place the cursor on any part of the performance trend to view the timestamp, the performance at that point in time, and the difference between the predicted performance and the threshold value defined for that resource. Because each target metric typically covers multiple resources (selected in the **Report Definitions**), you can isolate any given resource performance trend by selecting the resource from the list to the right of the graph. To check more resources, click **Select Component** and add target resources.

Sort by

The report is organized by target metric. View a target metric performance graph by clicking the metric on the left pane. You can select and view multiple graphs.

Sort the list of metrics on the left pane by selecting either **Status** (default) or **Metric**. **Status** organizes by severity of the potential alerts.

Additional actions

When you select resources from the resource list, you can make a note for further tracking. Additionally, click **Execute Action** to send an email message, run a script, or invoke the **Action Template** menu.

Report options

Select **Save the report as a Report Definition** to save the profile and target resource selections.

If conditions have changed or if you have run corrective actions in response to the initial performance trend projection, click **Re-run** to generate a new report using the same profile and report definitions. Then assess whether your response has averted the risk.

All generated risk reports are listed on the **Report Result** tab, where they can be sorted, viewed, or deleted.

Enhancing infrastructure management capabilities

You can analyze predicted data in the following contexts to understand performance trends in your infrastructure and allow for better decision making.

Analyzing near-term capacity trends

You can use predictive analytics risk reporting to analyze near-term capacity trends in your infrastructure. Infrastructure administrators know that keeping up with the growing capacity needs of application users is an ongoing effort. By using predictive analytics risk reporting, you can identify capacity needs on a weekly or daily basis by analyzing short-term trend projections. Use the following report definitions to analyze capacity usage trends:

- Top 5 Pools - Least Available Space
- Top 5 Pools Consuming Capacity (Free Space)
- Top 5 System Resources with the Growing Workloads
- Top 5 User Resources with the Growing Workloads

After you identify the affected storage resources, you can take the following actions:

- Assign more volumes manually.
- Create resource assignment rules based on your evaluation of the risk report data.
- Develop a plan for redistributing storage resources by moving storage pools to different volumes.

Long-term capacity planning

You can plan for future growth by estimating long-term capacity usage trend projections with predictive analytics risk reporting. Capacity planning poses several challenges for infrastructure administrators. Estimates for growth can fluctuate depending on the time of the estimate. When you use predictive analytics risk reporting, you can continually generate estimates for capacity growth. Begin by analyzing near-term trends. Use the following system-defined report definitions to collect data:

- Top 5 Pools Consuming Capacity
- Top 5 User Resources with the Busiest Workloads
- Top 5 System Resources with the Busiest Workloads
- Top 5 Growing Busy System Resources

Once you get an idea for the trend pattern for the various consumers or application users, create a risk profile that specifies capacity metrics. To maximize the information shown in the trend projection, create one profile for one month, one for three months, and one for six months. Use these profiles in conjunction with report definition that includes consumers.

Near-term performance tuning

You can use predictive analytics risk reporting to fine-tune overall system performance over time to achieve optimization. Because most large-scale IT infrastructures are heterogenous, the task of fine-tuning performance is ongoing. Infrastructure administrators might have the twin goals to make the most of the existing IT equipment (all system resources), while managing resources to achieve stability in performance. With the predictive analytics risk reporting feature in Ops Center Analyzer, you can generate reports with trend projections to anticipate performance fluctuations and adjust resource monitoring.

In this workflow, infrastructure administrators can fine-tune system resource performance and user resource allocation on a weekly or monthly basis.

1. Set up dynamic thresholds for user and system resources. Use the base dynamic threshold profile or edit it for the following:
 - Metrics: If you want to track performance at a granular level, select more metrics.
 - Plan: You can adjust the profile to monitor during peak times during the week and month
2. Set up event notifications.
3. Create risk profiles and report definitions to the corresponding metrics and resources in the profiles.
4. Generate a risk report along with standard reports to compare past and current trends with the trend projections in the risk report.
5. Make adjustments in resource allocation, and modify both threshold profiles as needed.
6. Generate the next round of risk reports and determine if the adjustments averted overutilization of resources.

As an ongoing process, infrastructure administrators have several options on how to act on the information in these risk reports. Ops Center Analyzer offers the following tools to aid in fine-tuning:

- You can develop a script that is invoked when certain performance events and threshold violations occur, and use the Execute Action function to run that script.
- You can use Execute Action to invoke a service from Ops Center Automator.
- If the magnitude of your infrastructure requires thorough analysis of data to avoid service interruptions, you can do the following:
 - To analyze minute-level data, adjust the collection time in Analyzer detail view to track events in finer granularity.
 - To analyze second-level data in Hitachi resources from Ops Center Analyzer, run granular data collection.

Damage Control

You can use the predictive analytics risk reporting feature to evaluate near-term trends and take preventive measures against performance degradation in your infrastructure.

Performance degradation can affect infrastructure with sudden bottlenecks. or worse, extended outages. Consumers might experience I/O problems with lagging response times, an annoying occurrence to application users, and a headache for IT. However, in an extended outage, the application is no longer available for use, causing a work stoppage. This situation can be more than a headache to IT and might result in escalating support calls.

In this workflow description, an infrastructure administrator uses predictive risk reporting in Ops Center Analyzer to respond preemptively to a sudden decrease in performance:

1. Monitor near-term trends across your infrastructure by using the following report definitions.
 - Top 5 Resources with Worst Response-times and Related Workloads
 - Top 5 Pools - Least Available Space

- Top 5 User Resources with the Highest Resources
 - Top 5 Platinum Consumers at Risk
2. When you isolate which resources consistently appear in these reports, create a new risk profile and risk report definition to analyze trend projections for those resources.
 3. Use the new profile and report definition to run risk reports projecting when the performance degradation will occur.
 4. Initiate immediate action to avert performance degradation:
 - Assign resources
 - Set up resource assignment rules
 5. If you determine that the performance trends are recurring, develop a response plan:
 - Adjust thresholds or create a new threshold profile for the resources
 - Set up notifications to track and alert other system administrators
 - Execute a script
 - Run Execute Action

Use risk reporting as an extra layer of monitoring

You can add an extra layer of resource monitoring to give a 10% margin outside your normal threshold limits during day-do-day operations by using predictive analytics risk reporting. To manage performance risks on a daily level, you can generate risk reports to establish a buffer for certain thresholds. Doing so allows infrastructure administrators to predict when thresholds violations will occur. This buffer makes it easier to react to performance problems.

Periodic analysis with predictive analytics risk reporting

You can generate predictive analytics risk reports for the same target resources and metrics to analyze near-term or long-term trends by using risk report definitions periodically.

The following workflow diagram illustrates how to use predictive analytics risk reporting for periodic analysis.

The image illustrates a three-step process for analyzing performance trends with predictive analytics risk reporting in Hitachi Ops Center Analyzer.

Step 1: The user navigates to the **Predictive Analytics** section and selects a report from the **Predict Result** list. The selected report is **Top 5 Resources with Slowest Response Times**.

Step 2: The user views the **Show Result** page, which displays a table of report execution history. The selected report is highlighted in blue.

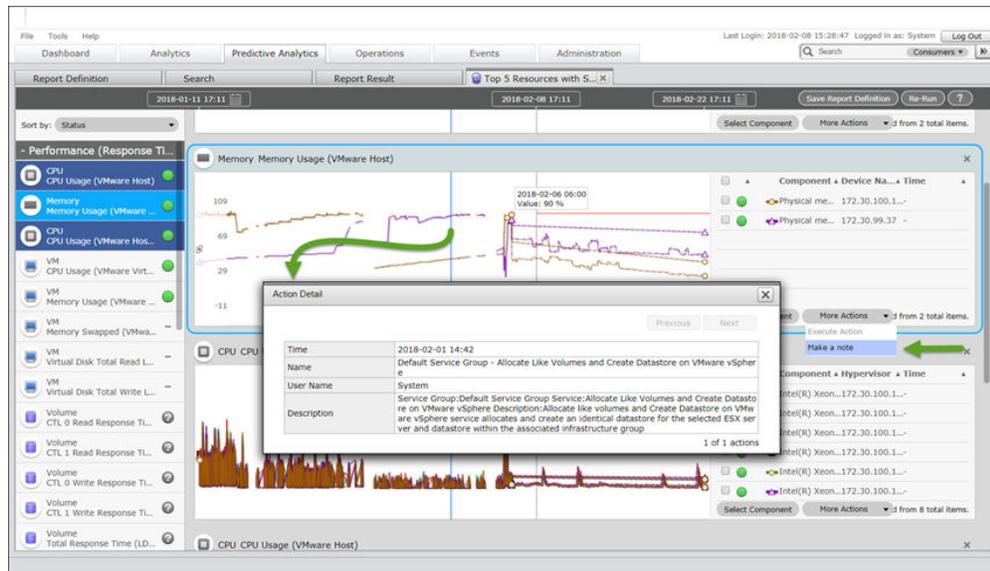
Name	Task Status	Report Status	Time	Time of Most Recent Action	Generated By
OrDemand_2018-01-31_09:16:34	Succeeded	Normal	2018-01-30 16:16:34	-	System
Demo_Report	Succeeded	Normal	2018-01-30 16:00:43	2018-01-19 16:22:00	System
Top 5 Busiest System Resources	Succeeded	Critical	2018-01-30 15:57:52	2018-01-20 00:52:00	System
OrDemand_2018-01-30_14:48:32	Failed	Unknown	2018-01-29 21:48:32	-	System
OrDemand_2018-01-30_14:48:13	Failed	Unknown	2018-01-29 21:48:13	-	System
Top 5 Pools - Least Available Space	Succeeded	Normal	2018-01-29 07:21:55	-	System
Top 5 Pools Consuming Capacity (Free Space)	Succeeded	Normal	2018-01-29 07:21:55	-	System
Top 5 Pools Consuming Capacity (Free Space)	Succeeded	Normal	2018-01-26 15:26:00	-	System
Demo_Report	Succeeded	Normal	2018-01-26 01:41:48	2018-01-19 16:22:00	System
OrDemand_2018-01-26_18:14:20	Succeeded	Normal	2018-01-26 01:14:20	-	System
Top 5 Busiest System Resources	Succeeded	Critical	2018-01-26 00:46:19	2018-01-20 00:52:00	System
Platinum Consumers at Risk	Succeeded	N/A	2018-01-25 23:57:28	-	System
Top 5 Resources with Slowest Response Times	Succeeded	Normal	2018-01-25 23:54:55	-	System
Demo_Report	Succeeded	Normal	2018-01-25 23:52:29	2018-01-19 16:22:00	System
Demo_Report	Succeeded	Normal	2018-01-25 23:50:47	2018-01-19 16:22:00	System
Top 5 Busiest System Resources	Succeeded	Critical	2018-01-25 17:09:47	2018-01-20 00:52:00	System
OrDemand_2018-01-26_09:21:07	Succeeded	Normal	2018-01-25 16:21:07	-	System
Top 5 Busiest System Resources	Succeeded	Critical	2018-01-25 04:01:03	2018-01-20 00:52:00	System
Top 5 Resources with Slowest Response Times	Succeeded	Normal	2018-01-25 01:22:13	-	System
Platinum Consumers at Risk	Succeeded	Critical	2018-01-25 01:20:06	-	System
Top 5 Busiest System Resources	Succeeded	Critical	2018-01-24 23:37:32	2018-01-20 00:52:00	System
OrDemand_2018-01-26_22:32:35	Succeeded	Normal	2018-01-24 23:22:35	-	System
Top 5 Busiest System Resources	Succeeded	Critical	2018-01-24 22:41:18	2018-01-20 00:52:00	System
OrDemand_2018-01-25_15:37:57	Succeeded	Normal	2018-01-24 22:37:57	-	System
OrDemand_2018-01-25_11:36:42	Succeeded	Normal	2018-01-24 22:36:42	-	System
Predictor_for_Systems_Vms	Succeeded	Critical	2018-01-24 22:25:46	-	System
OrDemand_2018-01-25_14:15:40	Failed	Unknown	2018-01-24 21:15:40	-	System
Top 5 Busiest System Resources-test 03	Succeeded	Critical	2018-01-23 20:36:10	-	System
Top 5 User Resources with the Growing Workloads	Succeeded	N/A	2018-01-18 10:42:58	-	System
Top 5 Resources with Worst Response Times & Related Wor...	Succeeded	Normal	2018-01-17 14:17:49	-	System

Step 3: The user analyzes the performance metrics for the selected report. The **Report Definition** is **Top 5 Resources with S... X**. The **Report Result** shows a value of 14.7976% on 2018-02-04 20:14, which is 75.202564% (Under) from the threshold of 75.202564% (Under). The predictive modeling type is 6% Regression. The performance metrics are visualized in three charts: **VM Memory Usage (VMware Virtual Machine)**, **CPU CPU Usage (VMware Host)**, and **VM Memory Usage (VMware Virtual Machine)** (repeated).

1. On Predictive Analytics > Report Definition, select a risk report definition, then click Predict Result.
2. On the Report Results tab, select the risk report and click Show Result.
3. The predictive analytics risk report opens in a new tab.

Recurring analysis

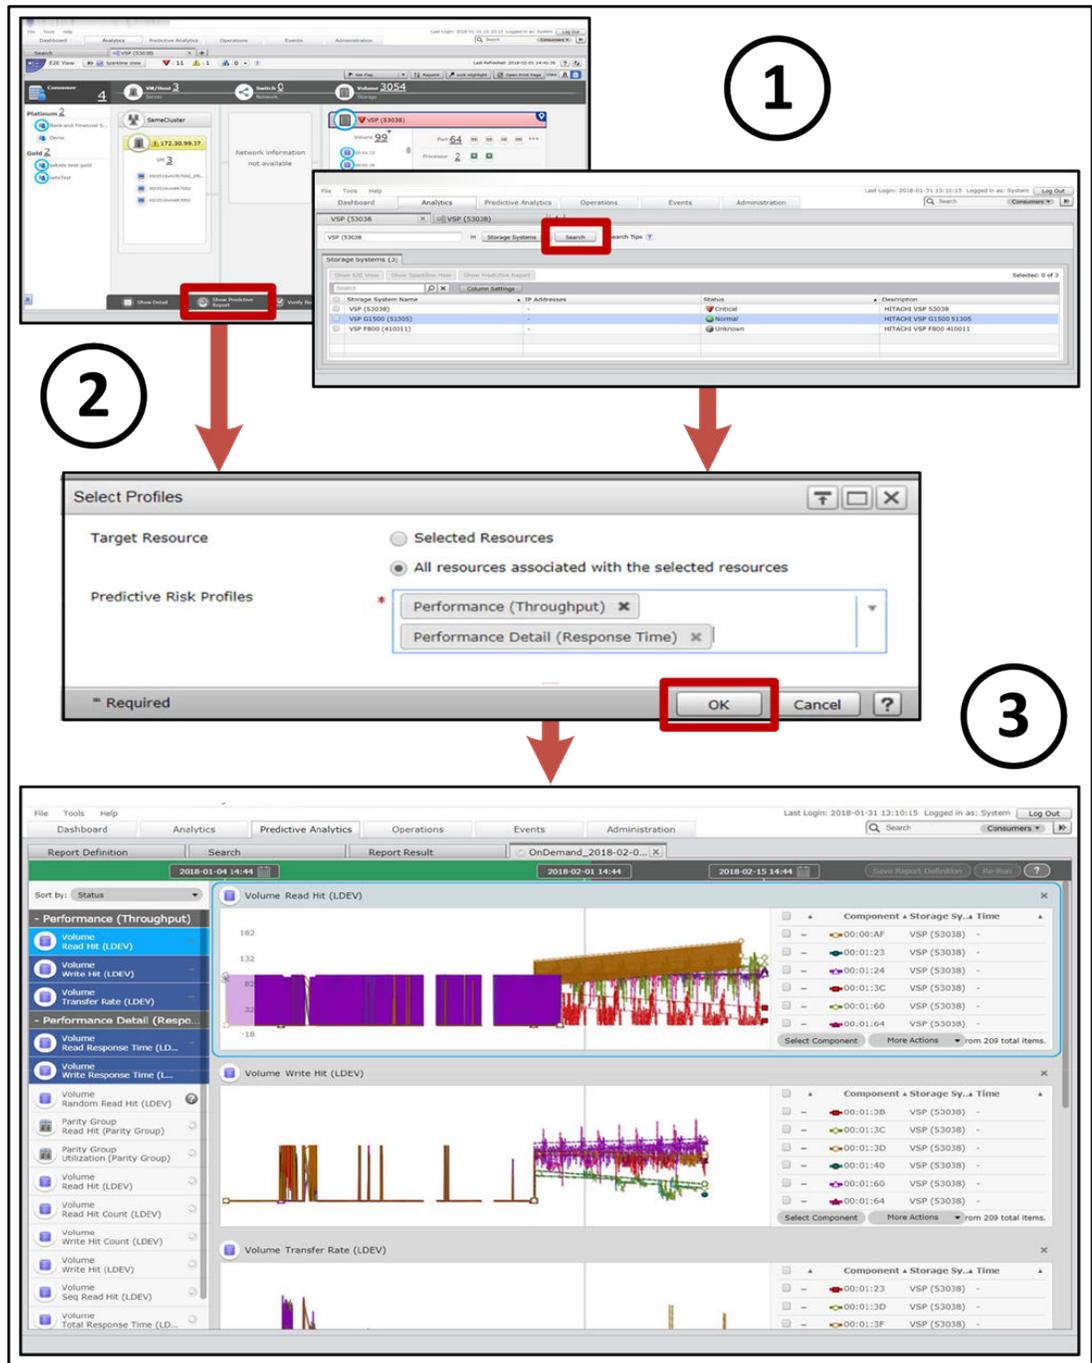
Place the cursor on any part of the performance trend to view the timestamp, the performance at that point in time, and the difference between the predicted performance and the threshold value defined for that resource. When you select resources from the resource list, you can make a note for further tracking.



E2E View analysis with predictive analytics risk reporting

You can generate predictive analytics risk reports from the **E2E View** tab or **Search** fields to examine the relationship of resources to the target resource based on current conditions.

The following workflow illustrates how to use predictive analytics risk reporting from the E2E View.



1. On the **E2E View** tab, select a resource for which you want to view predicted performance data.
2. When the Select Profiles dialog appears, select a **Target Resources** option and **Predictive Risk Profiles**, then click **OK**.
3. The predictive analytics risk report opens in a new tab.

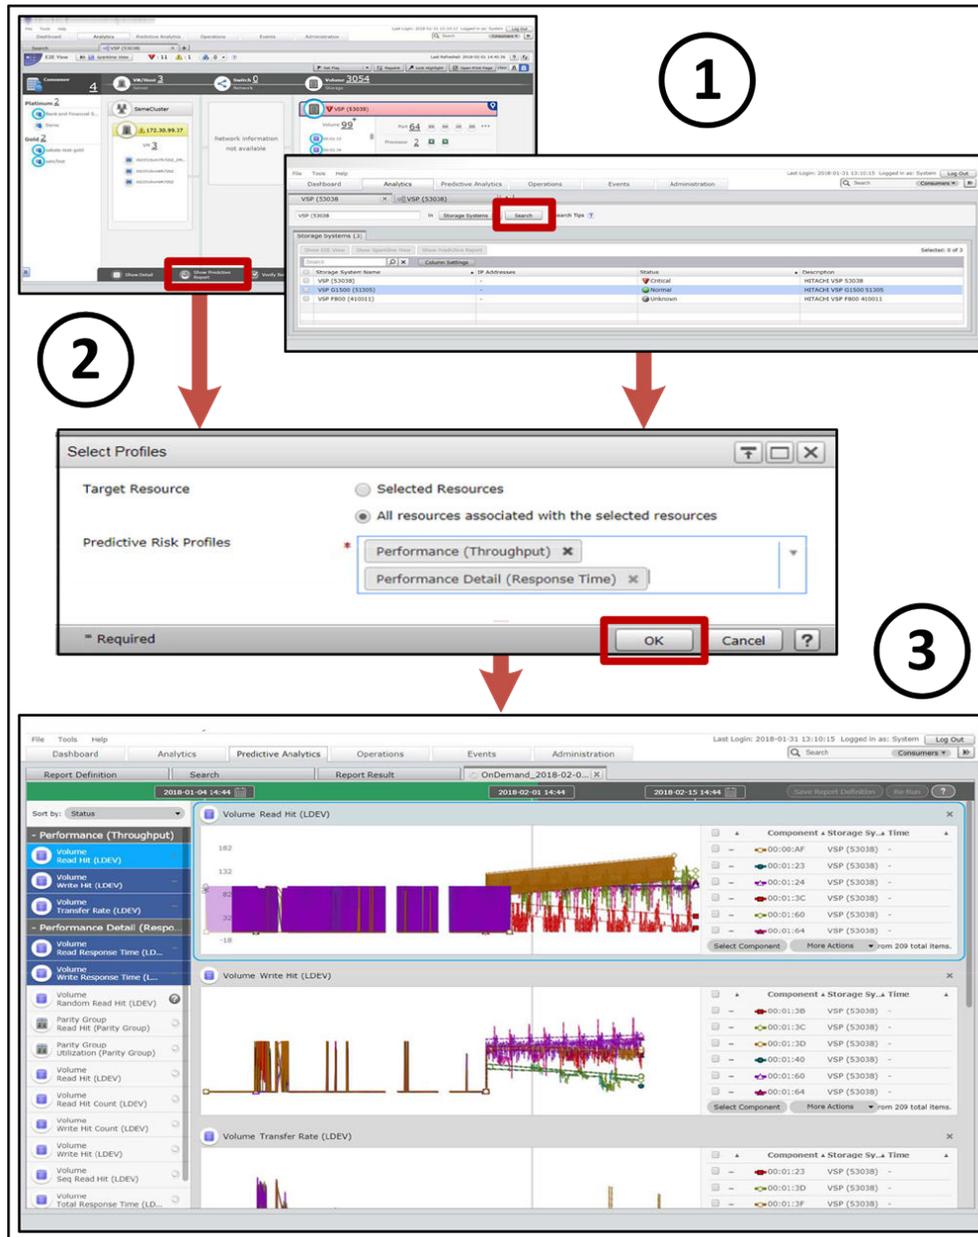
Target resource analysis

You can choose to analyze only selected resources from the E2E View or all resources associated with those selected resources.

Preventive actions and predictive analytics risk reporting

You can generate predictive analytics risk reports to check on the effects of corrective or preventive actions by comparing actual performance data with predicted data.

The following workflow diagram illustrates how to use predictive analytics risk reporting for follow-up analysis after taking corrective or preventive actions.



1. On Predictive Analytics > Report Definition, select a risk report definition, then click Predict Result.
2. On the Report Results tab, select the risk report and click Show Result.

3. The predictive analytics risk report opens in a new tab.
 - Select a resource, then Execute Action to do any of the following:
 - Send an email
 - Run a user-defined script
 - Start a service using Ops Center Automator
 - If conditions have changed or if you have run corrective actions in response to the initial performance trend projection, click Re-run to generate a new report, using the same profile and report definitions.

Chapter 6: Optimizing infrastructure resources with storage I/O controls

Ops Center Analyzer provides storage I/O controls to optimize infrastructure resources.

This feature works in many contexts to improve the efficient usage of resources in your infrastructure.

I/O control overview

I/O controls are available in Ops Center Analyzer when your storage systems have the Server Priority Manager function enabled.

Invoke this functionality through the Configuration Manager REST API and Hitachi Ops Center Automator, which runs the storage I/O control service. Configuration tasks include registering storage systems on the Configuration Manager REST API server and setting up a connection between Ops Center Analyzer and Ops Center Automator management servers. Alternately, if you do not have Ops Center Automator installed in your environment, you can develop a script to run the I/O control service. To use the script, specify the file path in the template files. Both tasks must be completed by the user.

I/O control operations

In Ops Center Analyzer, use I/O controls to set upper limits on either the I/Os per second or the data transfer rates across multiple storage systems at the volume level.

In the Ops Center Analyzer Operations window, search for and set limits on target volumes. The Operations window includes a History tab that displays a list of tasks created for each I/O control setting enabled in Ops Center Analyzer. This tab allows storage administrators to confirm that the upper limit setting is processed and to track the status of multiple tasks.

Although Server Priority Manager allows upper limits to be set on I/O activity at the port and WWN level, Ops Center Analyzer does not currently support these operations.



Note: Review and record any existing Server Priority Manager performance tuning settings on your storage systems before you set up I/O controls in Ops Center Analyzer. I/O upper limits set on volumes through Server Priority Manager are overwritten by upper limits set on Ops Center Analyzer. For details about Server Priority Manager, see the Performance Monitoring Guide available with your storage system.

Requirements

Ops Center Analyzer I/O control settings depend on the following requirements:

- The Configuration Manager REST API and Ops Center Automator must be installed.
- The target storage systems must have the Server Priority Manager function enabled.
- You must have an Administrator account on the storage system service processor (SVP) with Create permissions to connect the storage system and Ops Center Automator.
- You must register the target storage systems on the Configuration Manager REST API server.
- The StorageOps permissions must be set in Ops Center Analyzer to allow you to set, modify, or clear an upper limit on I/O traffic from the server to storage system.

Search capabilities for target volumes

Ops Center Analyzer provides search capabilities for finding target volumes on which you want to set upper limits.

In the Operations window, search for volumes by keyword using the following categories:

- Tasks: name or description of the task entered in the Set IO Control window.
- Consumers: consumer names, grades, or descriptions.
- Volumes: volume names, storage systems, or host groups.

If you leave the search field empty and click Search, the tab displays all consumers, volumes, and tasks. When selecting targets to set the upper limit, you can view all the associated volumes for the consumer or task. This feature is useful when you need to modify I/O control limits that have already been set.

Upper limit setting metrics

Choose a metric for which to set an upper limit.

When setting upper limits, the I/O control settings dialog gives you the option of applying one of the following metrics:

- IOPS: Use if the target volume is issuing a high number of I/O requests. The maximum value is 65,535 IOPS.
- Transfer Rate [MBps]: Use if the target consumer's I/O requests are small in number but large in data size. The maximum value is 31 MBps.

To confirm the result after setting upper limits, track the task status in the History tab in the Operations window.

Use cases for setting upper limits

Use I/O controls to optimize infrastructure resources.

Set upper limits for the following use cases:

- To achieve overall optimization of infrastructure resources during periods of I/O-intensive activity. In this use case, set temporary limits as a preemptive solution. Modify the upper limit according to known patterns in resource usage.
- To maintain a quality-of-service benchmark for a service-level objective (SLO). In this use case, set the upper limits as a long-term solution.
- To prioritize I/O activity for high-grade consumers after detecting bottlenecks that are affecting their performance. Disable the I/O control settings when the traffic between the server and storage system decreases to acceptable levels.

I/O control settings for an SLO

I/O controls enable you to meet the goals of your SLO.

SLAs specify a quality-of-service benchmark for an SLO. For storage I/O throughput, this benchmark is typically measured in IOPS or MBps. As a preemptive measure, Ops Center Analyzer enables you to set limits on storage I/O activity for applications on servers that issue too many I/O requests, and therefore provide sufficient resources in the infrastructure to meet the SLOs. After identifying the consumers with a specific SLO, select the volumes and set an upper limit to guarantee the quality-of-service benchmark for that SLO. You can set different storage I/O upper limits for consumers based on grade.

I/O controls for optimizing infrastructure resources

Ops Center Analyzer features I/O controls for planning the optimization of the resources in your infrastructure as usage patterns of critical and noncritical applications dictate.

Optimize your infrastructure resources by setting limits on I/O for noncritical applications. Setting upper limits on noncritical applications is similar to applying caps on I/O usage to free up more resources in the infrastructure. When you foresee increased I/O activity, set upper limits on volumes associated with applications or host servers issuing I/O requests.

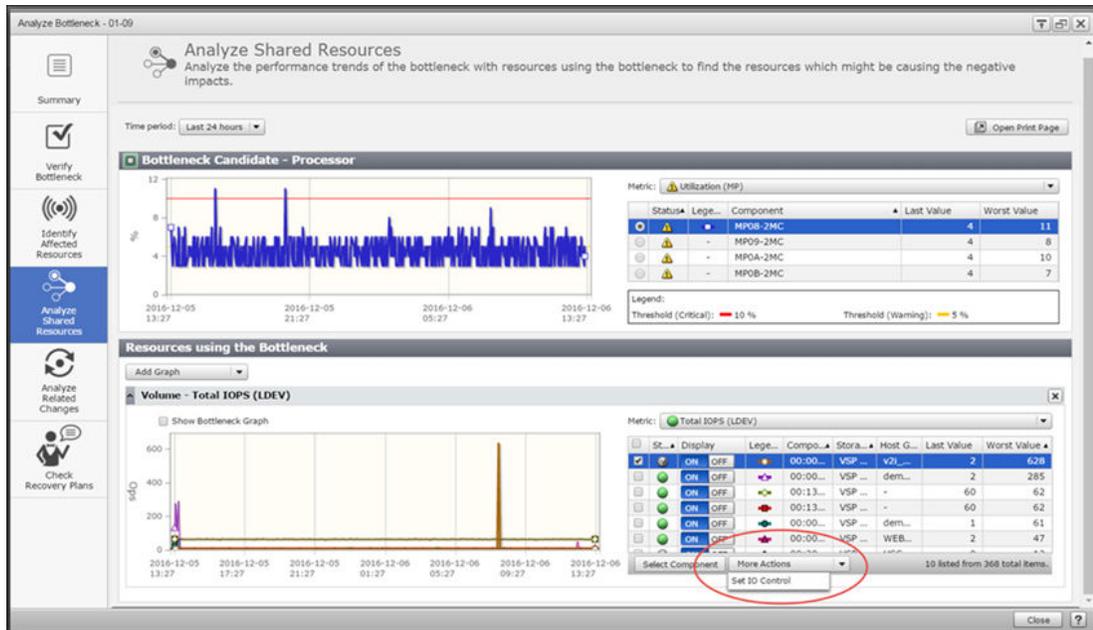
When development or testing efforts require more resources than usual, set I/O control limits on the volumes associated with the I/O-intensive applications for that period. This I/O control setting then allows the business-critical applications sufficient access to storage resources.

When I/O activity decreases to acceptable levels, clear the I/O control limits from those volumes. By establishing these temporary limits on storage I/O, the infrastructure achieves overall optimization during periods of increased I/O activity.

I/O controls for optimizing performance after the bottleneck analysis

To prevent an increased workload from affecting critical resources, set upper limits for servers issuing many I/O requests and affecting critical resources.

Storage administrators must respond quickly to sudden changes in I/O traffic. Shared infrastructure resources can degrade in performance at unpredictable times. If the bottleneck analysis detects a spike in total IOPS, as shown in the following figure, the root cause is an insufficient amount of resources available.



Because adding resources cannot be performed quickly, or might not be possible, the most efficient solution is to manage the I/O traffic. For a storage administrator, this situation must be treated as an emergency. In the Set IO Control window, set an upper limit for the volumes overusing resources immediately when you detect them.

You might use the upper limit setting as a temporary measure to manage the resources as planned for daily operations. If critical resources require less I/O prominence, remove the upper limit setting. All I/O control tasks are saved to the History tab.

Continue checking the **History** tab to monitor the upper limit settings by searching for tasks.

Preventing noncritical resources from causing performance degradation

When you are notified of performance degradation through an alert, perform the bottleneck analysis to detect the disruptive resource:

- Review the trend charts through E2E or Sparkline View to compare performance of selected resources.
- Use the Analyze Shared Resources window to identify the noncritical resources that are disrupting I/O traffic.
- When the Resources using the bottleneck window opens, you see a list of volumes that correspond to the trend chart.
- Identify the target volumes issuing many I/O requests.
- Select the target volume and then click More Actions to select Set IO Control. For your reference, give the task an appropriate name in the description field.
- Continue monitoring the History tab. When I/O control is no longer needed, select the target volumes of the task in IO Control Settings and click Off, or modify those limits as needed and resubmit.

Performing storage I/O control tasks in Ops Center Analyzer with Ops Center Automator

Use I/O controls to set, modify, or delete upper limits on I/O activity.

Setting I/O control limits

Set I/O control limits on volumes associated with a consumer. Use this procedure to set new or modify current I/O tasks.

Before you begin

- You must be logged on as a user with **StorageOps** permissions.
- The Server Priority Manager function must be enabled in the storage system.
- The connection between Ops Center Analyzer and Ops Center Automator is set with SSL communication.
- Check for existing Server Priority Manager settings in the I/O operations tab after you search for target volumes. You can see in the resulting list of volumes existing Server Priority Manager upper limit settings. If another administrator has set upper limits through Server Priority Manager at the volume level, running I/O control tasks in Ops Center Analyzer will overwrite the Server Priority Manager settings on the target volumes. However, if an upper limit has been set in Server Priority Manager at the port or WWN level, the I/O control task submitted on Ops Center Analyzer will not overwrite the previous Server Priority Manager upper limit setting.
- Review the list of upper limit tasks in the Operations > History > History tab, to verify whether any active tasks are in progress for the same target volumes for which you want to set I/O controls. If there are any tasks with the status (*In Progress*), wait for the task to complete (*Complete*) before you set new I/O control settings.
- The maximum number of WWNs associated with a target volume is 32. You cannot set upper limits on volumes that exceed this number.

Procedure

1. From the **Operations** tab, search for and select the volumes, and click **Set IO Control**. Alternatively, if you are troubleshooting a performance bottleneck in the **Analyze Shared Resources** window, select the target volumes, and click **Set IO Control**.



Note: If you are modifying the **Target Volumes** list, click **Remove Volumes**. A window opens where you can add or remove volumes from the **Set IO Control** task.

2. In the **Set IO Control** window, make the following selections:
 - a. In **Upper Limit Setting**, select **ON**.
 - b. In **Collective Settings**, select the metric and enter the limit in **Upper limit for each volume**.

- c. Enter a task name and description, and then click **Next**.

A default task name based on the date and time is automatically assigned:
 yyyyymmdd_hhmm_IOControlSettings.

3. In **Confirm the settings**, review the information and click **Submit**. Confirm the status of the task in the **History** tab by selecting **View task status**.



Tip: Ops Center Analyzer does not automatically show the progress of the I/O control task. If the I/O control task does not appear on the **History** tab, click **Update Status**.

Task Name	Registered Time of Task	User Name	Status	Last Status update	Description
20160916_1427_IOControlSettings	2016-09-16 02:07:49	System	Completed	2016-09-16 02:07:47	-
name.name	2016-09-16 02:18:37	System	Completed	2016-09-16 02:35:28	-
name.name2	2016-09-16 02:19:54	System	Completed	2016-09-16 02:20:49	-
sakita_20160916_1854_IOCentr	2016-09-16 02:54:16	System	In Progress	2016-09-16 02:54:19	-
sakita_delete_20160916_2044_I	2016-09-16 04:45:07	System	In Progress	2016-09-16 04:45:10	-
task.name	2016-09-16 02:07:24	System	Completed	2016-09-16 02:07:47	desc

In Progress

The I/O control task is submitted, but the upper limit setting is not in effect.

Completed

The I/O control tasks is complete, and the upper limit setting is in effect.



Note: If you submit the I/O control task through a script, the status is **Executed** instead of **Completed**.

Not Completed

The I/O control task is cancelled. The upper limit is not set.



Note: When you delete the volume of a storage system, the volume does not display in the I/O control task details.

Clearing I/O control limits

Clear I/O control limits when there is a change in storage I/O priorities.

Procedure

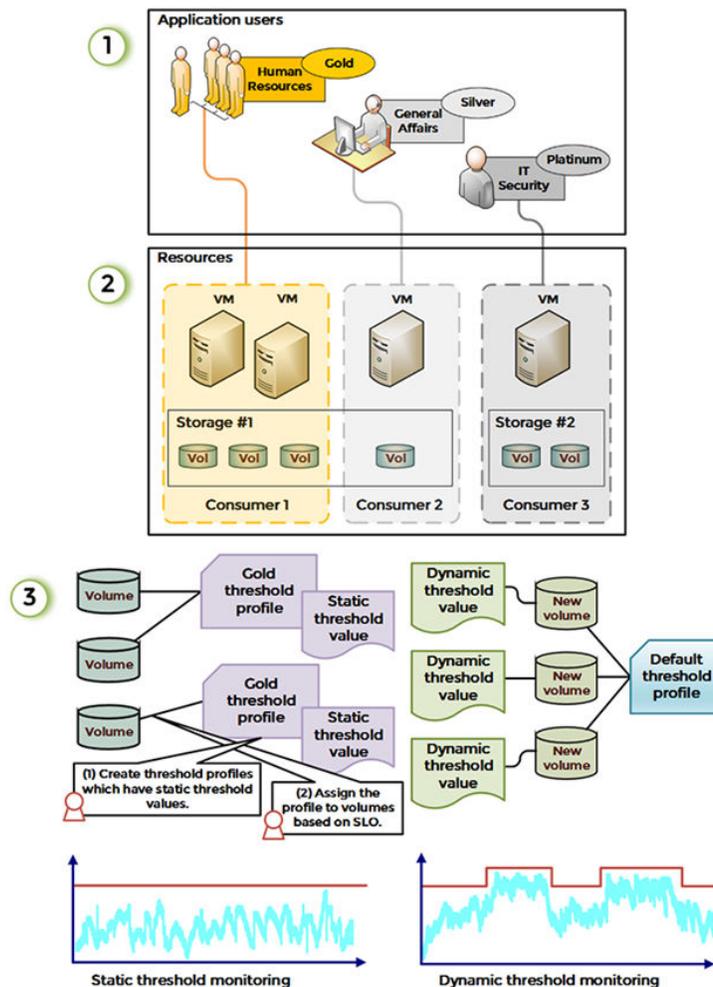
1. From **Operations**, search for volumes on which I/O control limits have been enabled.
2. Click **Set IO Control**, and a window opens.
 - a. In **Upper Limit Setting**, select **Off**.
 - b. Click **Next** and verify the information.
3. Click **Submit**, and then go to the **History** tab to verify the status of the new task.

Chapter 7: Configuring resource monitoring

Ops Center Analyzer ensures the health of your data center by measuring, monitoring, and optimizing the performance of your infrastructure resources.

Overview of managing and monitoring infrastructure resources

You can set up consumers associated with user groups or applications. You can then assign resources to those consumers and set up threshold profiles for monitoring those resources.



1. Create a consumer profile and set the grade: [Creating a consumer \(on page 110\)](#).

2. Add resources to the consumer automatically: [Creating a resource assignment rule \(on page 110\)](#).
3. Create threshold profiles to manage monitored resources: [Setting thresholds for system resources \(on page 105\)](#).

You can choose from two types of resource monitoring:

- For monitoring SLOs, see [Monitoring using a static threshold \(on page 99\)](#).
- For stability focused monitoring, see [Monitoring using a dynamic threshold \(on page 94\)](#).

You can also add resources either manually or configure user resource assignment rules based on threshold profiles: [Creating a user resource assignment rule \(on page 104\)](#).

Resource monitoring settings

To monitor management targets, define the following conditions:

- User Resource Threshold Profiles

Define the monitoring conditions for detecting deterioration in the service performance of virtual machines and volumes. You can monitor using dynamic thresholds or static thresholds.

You can create rules and conditions to automate resource assignment to monitoring profiles. Using these rules, the newly discovered user resources are automatically assigned to the user resource threshold profiles. When you do not create monitoring threshold profiles or define assignment rules, the newly discovered resources are automatically registered to the default threshold profiles.

- System Resource Threshold Profiles

Define monitoring conditions for detecting performance bottlenecks in infrastructure devices such as switches, hypervisors and storage systems. You can monitor using static thresholds.

For monitoring profiles, the following conditions occur:

- When static thresholds are used, the measured values collected from the monitoring target are evaluated based on whether they are greater than or equal to the warning threshold.

Result: An event occurs when the measured value of the metric reaches the warning threshold.

- When you use static or dynamic thresholds, measured values from the monitoring target are evaluated based on whether they are greater than or equal to the critical threshold.

Result: An event occurs when the measured value of the metric reaches the critical threshold.

When you restart the Ops Center Analyzer service, the status of monitored resources can be delayed for 5 minutes or longer. During this period, the status displays as **Unknown**.

Dynamic thresholds

Dynamic thresholds are calculated automatically by analyzing the load pattern from the historical data. These values are adaptive in nature and change over a period of time depending on the performance of your resources, workload changes, and so on. You use dynamic thresholds to monitor the user resources, such as volumes, VMs, and hosts.

The scenarios for using dynamic thresholds to monitor your environment are as follows:

- When SLOs and other performance parameters are not established with the customer
- When you want to monitor your environment for stable performance and detect irregular behavior

For details about dynamic threshold monitoring, see [Dynamic Threshold Storage Resource Monitoring With Performance Analytics](#).

Advantages of dynamic thresholds

With changing business requirements and performance goals, monitoring performance of your environment using predefined static thresholds might not be a feasible solution. The static values are calculated through trial and error, which is often time-consuming. These values become out of context in the long-term and the settings must be re-evaluated to ensure compliance.

By automating the threshold setting you gain better visibility into your environment and performance trend patterns. Dynamic thresholds adapt to your environment and proactively sends alerts before the performance bottleneck occurs.

For more information about dynamic threshold monitoring, see [Dynamic Threshold Storage Resource Monitoring With Performance Analytics](#).

Determining if the computed value is correct

If the computed values match your requirements, you can continue to use the dynamic thresholds for monitoring your environment. If you receive too many false alerts, you can manually edit the dynamic threshold values. For example, during migration process, a resource might temporarily have a large number of disk I/Os, and you might receive a number of false alerts. In this situation, you can manually edit the baseline value to account for the temporary increase in the load, and then allow the system to dynamically adjust the baseline values when the stable operation is restored.

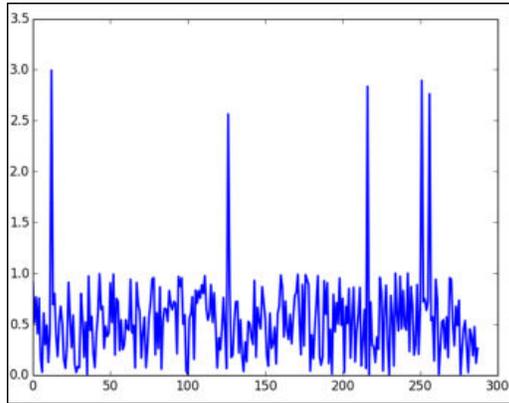
Automatic calculation of baseline values

Determining an appropriate threshold is essential while monitoring business-critical applications. Ops Center Analyzer analyzes the peak, normal, and low volume phases based on the historical data and adjusts the monitoring thresholds accordingly. Automating the threshold calculation eliminates false alerts and reduces the number of alerts to investigate, which might otherwise become a management overhead.

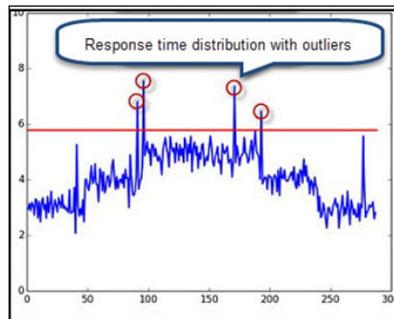
The application workloads might vary at different times of the day or week. For example, the workload pattern of an OLTP application might be different on weekdays and weekends. You can manage varying workloads that occur at different time periods for an application by creating monitoring plans. The system analyzes the performance data accumulated in the scheduled baseline period for computing the dynamic threshold values.

The following example shows the response time metrics of a business-critical application monitored over time and how the system derives the automatic threshold values based on the past performance. The high-level steps the system uses to calculate the automatic baseline values are as follows:

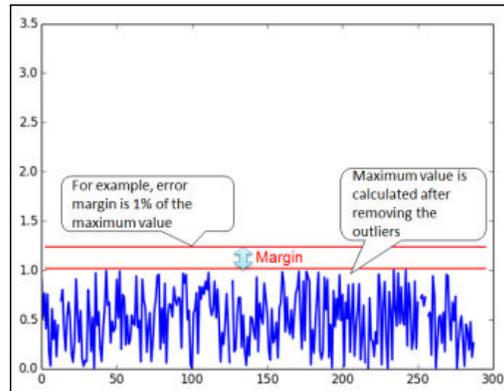
- Analyzes historical data for identifying the performance patterns in the specified baseline period.



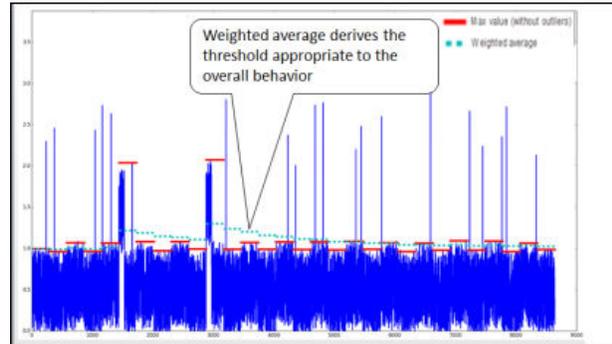
- Detects and removes the occasional outliers: In the following example, the data points that deviate from the norm represent the outliers. The system ignores the outliers appearing at irregular intervals to calculate an appropriate threshold value.



- Calculates the maximum value: The upper limit of the values in the normal range is used to calculate the maximum value. After determining the maximum value, the system adds the margin of error to the computed value.



- Determines the weighted average: The weighted average derives the threshold values based on the past performance trends over a specified time period.



Monitoring using a dynamic threshold

When monitoring using dynamic thresholds, the thresholds are automatically calculated and set based on performance information collected from monitored user resources over a specific time period. This type of monitoring is useful when you want to detect degradation in service performance or when, for example, you have not established a service-level agreement with customers. The appropriate thresholds are automatically set according to the resource configuration and changes in the load status. For this reason, system administrators do not need to manage the threshold values. In addition, based on the importance of the applicable system, you can set the margin levels for the calculated thresholds.

To monitor using dynamic thresholds, create threshold profiles for user resources and then assign the resources you want to monitor. You can set more than one monitoring plan for each user resource threshold profile. For the monitoring plan, specify threshold values for the monitoring schedule or for each monitoring item. By creating multiple monitoring plans, you can schedule changes to the threshold values or items to be monitored.

By creating assignment rules for user resources, the discovered user resources are automatically assigned to threshold profiles in accordance with the assignment rules. You can also edit and delete the created threshold profiles or assignment rules for user resources.

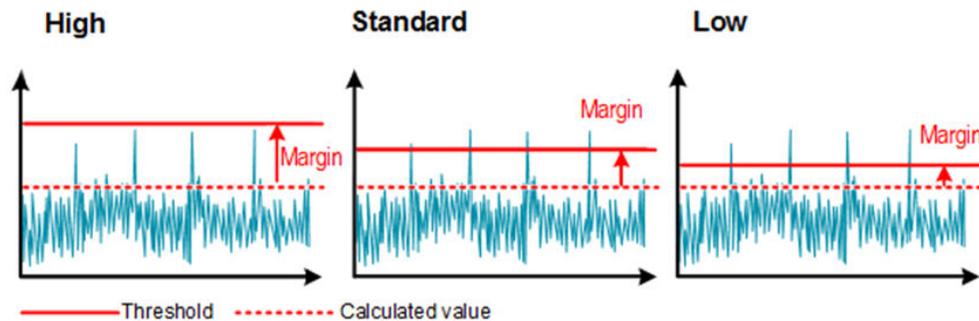
Note: Changing the monitoring mode of a profile to Dynamic or Off may result in the `Threshold: Could not get the value` error message when a performance event of a resource monitored by that profile occurs.

Dynamic threshold monitoring margins

Dynamic thresholds are based on computed values, allowing system administrators to detect performance degradation and take preemptive action.

This setting only affects the threshold profile with a dynamic monitoring mode. The threshold is automatically calculated based on the performance information, and a safety margin is added to the calculated value to avoid false alerts. The higher the safety margin, the lower the total number of alerts generated. Likewise, the lower the safety margin, the higher the total number of alerts.

The following graphs illustrate the differences between the three type of margins.



Selecting the dynamic threshold margin

Select the margin for dynamic thresholds and create a monitoring plan when creating a threshold profile.

Dynamic thresholds are automatically calculated based on performance information.

Choose the margin type based on your monitoring objectives:

- Standard: This is the recommended margin.
- High: Use this margin to generate a lower number of alerts.
- Low: Use this margin to generate a higher number of alerts.

You can change the margin on the dynamic threshold profile as performance trends change.

You can also create a monitoring plan if your environment requires flexible monitoring.

Enter the threshold settings for a specific resource type.

Profile Name: ODB1

Description:

Resource Type: Volume

Copy From: #Preset Profile for Volume

Margin for Dynamic Threshold: Standard

Pattern	Plan Name	Schedule	From	To	Priority	Description
	ODB_DTM1	Tue,Wed,Thu	06:00	18:00	1	Midweek
	Base	Mon,Tue,Wed,Thu,Fri,Sat,Sun	00:00	24:00	2	Base Plan

Monitoring Mode	Metric	Threshold	
		Warning	Critical
Dynamic	Total IOPS (LDEV)	-	Thresholds vary among resources
OFF	Read Hit (LDEV)	-	-
Dynamic	Read Hit Count (LDEV)	-	Thresholds vary among resources
OFF	Write Hit Count (LDEV)	-	-
Dynamic	Write Hit (LDEV)	-	Thresholds vary among resources
OFF	Read IOPS (LDEV)	-	-

* Required

OK Cancel ?

Create a plan by choosing a schedule and setting threshold values for various metrics:

- Select a resource that you plan to monitor.
- Devise a schedule and set threshold values for each metric.
- Select the dynamic threshold margin type.

The dynamic threshold values depend on the performance data of the monitored resources.

Static thresholds

Static thresholds are user-defined thresholds that you manually configure for use at different times of the day or week depending on the workload in your environment.

You use predefined static threshold values in the following scenarios:

- You have a well-defined service-level objective which clearly establishes the performance goals.

For example, if you have a service-level agreement with the customer to support online transactions at a response time of less than 1 second for a business-critical application, then you can create a User resource threshold profile to establish the response time and other performance requirements for the application, and then assign the target resources for monitoring. If there is an SLO violation, the system sends a critical alert or a warning and notifies the user before the problem becomes serious. You can also generate a report that compares the actual response time of the business-critical application to the SLO and see if your objectives are in compliance, then take necessary measures to fix the problem.

- You can assess the workload patterns in your environment and know what values to assign.

For example, define the threshold for a system resource based on the architecture of the storage system. If the storage system is VSP G1000, then the recommended MPB (MP Blade) usage is under 60%.

Setting static thresholds using monitoring profiles

You can create monitoring profiles with static thresholds for managing user and system resources. The performance parameters defined in the threshold profile determine when an alert is triggered.

Create threshold profiles for user or system resources based on the resource type, and then assign the resources you want to monitor.

For user resources

The procedure for setting static thresholds for user resources is as follows:

Procedure

1. On the **Administration** tab, from the navigation pane, select **Monitoring Settings > User Resource Threshold Profiles > Create Threshold Profile**.
2. In the **Create User Resource Threshold Profile** window, enter the profile name, description, and select the resource type.
3. On the **Monitoring Plans** tab, click **Create Plan** to create new monitoring plans. You can either edit the base plan or create a new plan.
4. In the **Create Plan** window, set the target period for monitoring. Under Target metric, click **Static** to enable static monitoring mode. You must manually enter the threshold values for the target metrics when you enable static monitoring mode.

Enter the time period, monitoring mode, and threshold values for the target metrics. You can create up to ten plans per profile.

Plan Name:

Description:

Copy From:

Period:

Target Days: Mon Tue Wed Thu Fri Sat Sun

Target Period: To (UTC - 6:00)

Target Metrics:

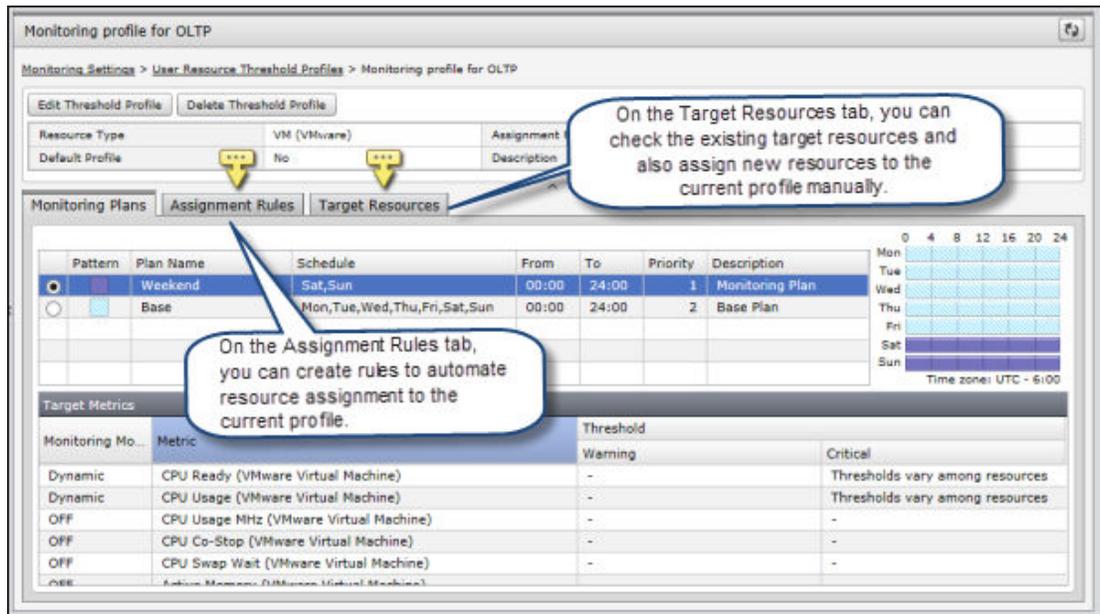
Monitoring Mode	Metric	Threshold							
		Warning	Critical						
Dynamic	Static	OFF	CPU Ready (VMware Virtu...)	>	<input type="text" value=""/>	msec	>	<input type="text" value=""/>	msec
Dynamic	Static	OFF	Virtual Disk Total Read La...	>	<input type="text" value=""/>	msec	>	<input type="text" value=""/>	msec
Dynamic	Static	OFF	Virtual Disk Total Write La...	>	<input type="text" value=""/>	msec	>	<input type="text" value=""/>	msec
Dynamic	Static	OFF	Dropped Rx (VMware Virt...	>	<input type="text" value="1"/>	Numb...	>	<input type="text" value="5"/>	Number
Dynamic	Static	OFF	Dropped Tx (VMware Virt...	>	<input type="text" value="1"/>	Numb...	>	<input type="text" value="5"/>	Number
Dynamic	Static	OFF	CPU Usage (VMware Virtu...	-			-		
Dynamic	Static	OFF	CPU Usage MHz (VMware...	-			-		
Dynamic	Static	OFF	CPU Co-Stop (VMware Vir...	-			-		
Dynamic	Static	OFF	CPU Swap Wait (VMware...	-			-		
Dynamic	Static	OFF	Active Memory (VMware...	-			-		
Dynamic	Static	OFF	Memory Consumed (VMw...	-			-		
Dynamic	Static	OFF	Memory Overhead (VMwa...	-			-		

* Required

OK Cancel

5. To save the profile, click **OK**.
After you save the profile, you are navigated to the profile detail window, where you can assign target resources, or create resource assignment rules.

6. In the profile detail window you can do the following:
 - On the **Assignment Rules** tab, you can create rules for assigning resources to the monitoring profile automatically.
 - On the **Target Resources** tab, you can assign the resources to the profile manually. You can also view the existing target resources associated with the monitoring profile.

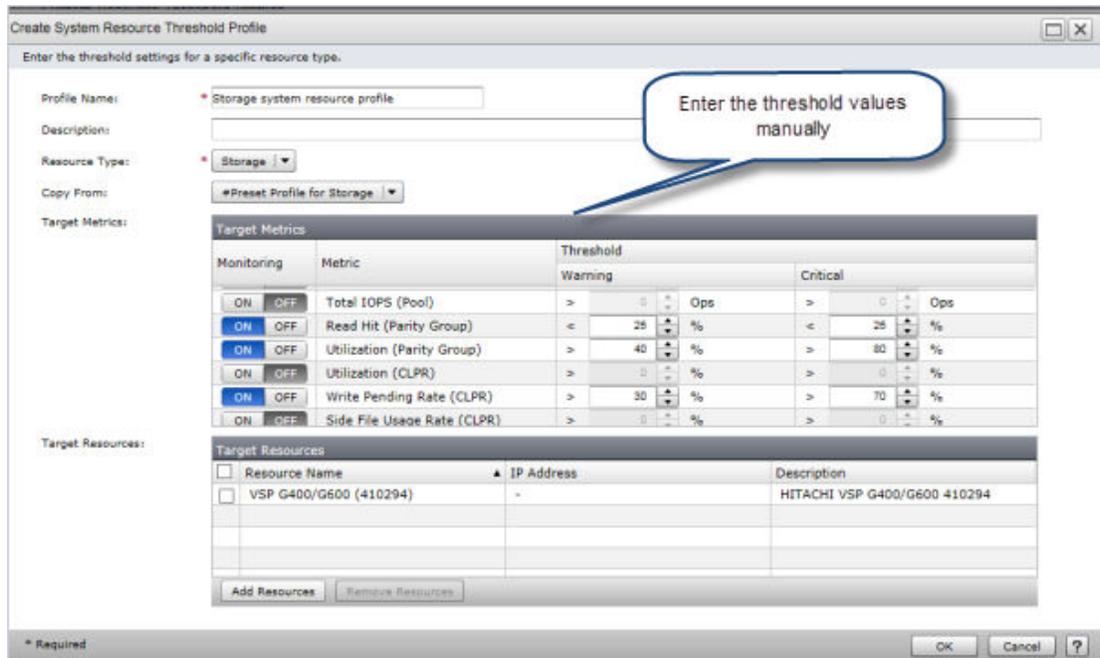


For system resources

The procedure for setting static thresholds for system resources is as follows:

Procedure

1. On the **Administration** tab, from the navigation pane select **Monitoring settings > System Resource Threshold Profiles > Create Threshold Profile**.
2. In the **Create System Resource Threshold Profile** window, enter the profile name, description, and select the resource type. If required, copy the settings from the default profile or an existing system resource profile.
3. Under threshold values, manually enter the threshold values for the performance metrics.



4. Under **Target Resources**, click **Add Resources** to manually assign resources to the system resource threshold profile.

Monitoring using a static threshold

When monitoring using static thresholds, you can detect service performance degradation and performance bottlenecks for user resources and system resources by setting thresholds for each monitored item based on the type of resource. This type of monitoring is useful when you have established a service-level agreement with users and threshold values can be defined based on the agreement, or when threshold values are defined based on the configuration of resources. You can set two threshold levels such as Critical and Warning. If the system detects that a threshold is exceeded, you are notified of the event. You can then verify the event details, including the names of the resources for which errors or warnings occurred.

To monitor using static thresholds, create threshold profiles for user resources or system resources based on the type of resources, and then assign the resources you want to monitor. You can set more than one monitoring plan for each user resource threshold profile. For the monitoring plan, specify threshold values for the monitoring schedule or for each monitoring item. By creating multiple monitoring plans, you can schedule changes to the threshold values or items to be monitored.

By creating assignment rules for user resources, the discovered resources are automatically assigned to threshold profiles in accordance with the assignment rules. You can also edit and delete the created threshold profiles or assignment rules.

About default profiles for volumes (user) and arrays (system)

Default monitoring profiles are provided at the user (hosts, VMs, and volumes) and system level (datastore, switches, hypervisors, and storage).

The User and System Resource Threshold Profiles are located in the Administration tab under Monitoring Settings. Each profile includes a list of metrics and values. The defaults for volumes (user) and arrays (system) are described here for reference. The User System Resource Threshold Profiles support [Dynamic thresholds \(on page 92\)](#). The default values are general recommendations and should be customized based on your environment.



Note: For best results, do not modify the default profiles; make a copy instead. (When creating a new threshold profile, use the Copy From option.)

Default Profile for Volumes

The default user profile for volumes are listed in the following table. The metrics can be set to Dynamic, Static, or OFF. Dynamic means the thresholds are calculated automatically based on historical data. Metrics that are set to Dynamic or OFF have no set values, hence the table does not include Threshold values.

Monitoring (OFF, Static/ Dynamic)	Metric
Dynamic	Total IOPS (Port)
OFF	Read Hit (LDEV)
OFF	Read Hit Count (LDEV)
OFF	Write Hit Count (LDEV)
OFF	Write Hit (LDEV)
OFF	Read IOPS (LDEV)
OFF	Random Read Hit (LDEV)
OFF	Random Read IOPS (LDEV)
OFF	Random Write IOPS (LDEV)
OFF	Seq Read Hit (LDEV)
OFF	Seq Read IOPS (LDEV)
OFF	Seq Write IOPS (LDEV)
OFF	Write IOPS (LDEV)
OFF	Backend Transfers (LDEV)
OFF	Cache to Drive Transfers (LDEV)
OFF	Random Drive to Cache Transfers (LDEV)
OFF	Seq Drive to Cache Transfers (LDEV)
OFF	Read Transfer Rate (LDEV)

Monitoring (OFF, Static/ Dynamic)	Metric
Dynamic	Transfer Rate (LDEV)
OFF	Write Transfer Rate (LDEV)
OFF	Random Read Transfer Rate (LDEV)
OFF	Random Write Transfer Rate (LDEV)
OFF	Seq Read Transfer Rate (LDEV)
OFF	Seq Write Transfer Rate (LDEV)
OFF	Read Response Time (LDEV)
Dynamic	Total Response Time (LDEV)
OFF	Write Response Time (LDEV)
OFF	Utilization (LDEV)

Default Profile for Storage

The default system profile for arrays includes a series of metrics with threshold values listed in the following table.

Monitoring (ON/OFF)	Metric	Threshold	
		Warning	Critical
OFF	Total IOPS (Port)	-	-
OFF	Max IOPS (Port)	-	-
OFF	Transfer Rate (Port)	-	-
OFF	Max Transfer Rate (Port)	-	-
OFF	Total IOPS (Pool)	-	-
ON	Usage Rate (Pool)	> 80 %	> 90 %
ON	Read Hit (Parity Group)	< 25 %	< 25 %
ON	Utilization (Parity Group)	> 40 %	> 80 %
ON	Write Pending Rate (CLPR)	> 30 %	> 70 %
ON	Utilization (Parity Group)	> 40 %	> 80 %
OFF	Utilization (CLPR)	-	-

Monitoring (ON/OFF)	Metric	Threshold	
		Warning	Critical
ON	Write Pending Rate (CLPR)	> 30 %	> 70 %
OFF	Side File Usage rate (CLPR)	-	-
ON	Utilization (MP)	> 40 %	> 80 %
ON	Write Pending Rate (MPB CLPR)	> 30 %	> 60 %
ON	Utilization (MPB)	> 40 %	> 80 %
ON	Access Path Usage (CHA ESW)	> 20 %	> 50 %
ON	Access Path Usage (DKA ESW)	> 20 %	> 50 %
ON	Access Path Usage (MPB ESW)	> 20 %	> 50 %
ON	Access Path Usage (Cache ESW)	> 20 %	> 50 %
ON	Access Path Usage (Cache Path)	> 20 %	> 50 %
ON	Physical Capacity Usage Rate (Pool)	> 80 %	> 90 %

Default Profile for VSSB

The default system profile for Virtual Storage Software Block includes a series of metrics with threshold values listed in the following table.

Monitoring (ON/OFF)	Metric	Threshold	
		Warning	Critical
ON	Average CPU Usage (VSSB Cluster)	> 40 %	> 60 %
OFF	Average Read Transfer Rate (VSSB Cluster)	-	-
OFF	Average Write Transfer Rate (VSSB Cluster)	-	-
ON	CPU Usage (VSSB CPU)	> 40 %	> 60 %

Monitoring (ON/OFF)	Metric	Threshold	
		Warning	Critical
OFF	Read IOPS (VSSB Compute Port)	-	-
OFF	Write IOPS (VSSB Compute Port)	-	-
ON	Read Transfer Rate (VSSB Compute Port)*	> 400 MBps	> 1000 MBps
ON	Write Transfer Rate (VSSB Compute Port)*	> 400 MBps	> 1000 MBps
OFF	Response Time (VSSB Compute Port)	-	-
ON	Receive Transfer Rate (VSSB Intermode Port)	> 400 MBps	> 800 MBps
ON	Send Transfer Rate (VSSB Intermode Port)	> 400 MBps	> 800 MBps
ON	Usage (VSSB Drive)	> 90 %	> 99 %
ON	Response Time (VSSB Drive)	> 0.5 msec	> 1 msec
OFF	Read IOPS (VSSB Drive)	-	-
OFF	Write IOPS (VSSB Drive)	-	-
OFF	Read Transfer rate (VSSB Drive)	-	-
OFF	Write Transfer Rate (VSSB Drive)	-	-
<p>Notes:</p> <p>* Set the threshold values according to your network environment:</p> <ul style="list-style-type: none"> ▪ 10 Gbps ISCSI - Warning: 400MiB/s, Critical: 1000MiB/s ▪ 25Gbps ISCSI - Warning:1000MiB/s, Critical: 2100MiB/s ▪ FC - Warning: 40% of the Port speed, Critical: 80% of the Port speed 			

Setting thresholds for user resources

To detect service performance degradation of a virtual machine or volume, create a profile and assign rules.

Procedure

1. In the **Administration** tab, select **Monitoring Settings > User Resource Threshold Profiles**.
2. Click **Create Threshold Profile**.
3. In the **Create User Resource Threshold Profile** window, define the Profile Name and Resource Type.
4. Select a resource to be monitored from **Resource Type**, and then click **Create Plan**.
5. In the displayed box, create a monitoring plan.
Define the necessary items, and then click **OK**.
6. For **Target Metric**, select a monitoring mode for each monitoring item. If you select **Static**, define a threshold. If you select **Dynamic**, go to the next step.
7. In the **Create User Resource Threshold Profile** window, click **OK**.
The window automatically switches to the detailed window for the user resource threshold profile that was created.
8. Click the **Assignment Rules** tab to create resource assignment rules.
9. Click **Create Rule**, and then specify the necessary items in the displayed box.
10. Define conditions for **Condition**, and then click **Preview**.
A list appears with the resources that match the specified conditions and are not assigned to the threshold profiles of other user resources.
11. Click **OK**.
The user resource threshold profile is associated with the resource, and thresholds are set for the resource.
Verify which resources are associated with a user resource threshold profile by clicking the **Target Resources** tab.

Result

The created profile appears in the user resource threshold profiles list.

Creating a user resource assignment rule

When you create a user resource assignment rules, the discovered user resources are automatically registered in a threshold profile 24 hours after discovery.

Procedure

1. In the **Administration** tab, select **Monitoring Settings > User Resource Assignment Rules**.
2. Select the tab for the user resource you are creating, and then click **Create Rule**.
3. Define the assignment conditions and the assignment destination.
4. Click **OK**.

Result

The created user resource assignment rule appears in the user resource assignment rules list.

Changing user resource assignment rule priority

Change the order of priority for the user resource assignment rule.

Procedure

1. In the **Administration** tab, select **Monitoring Settings > User Resource Assignment Rules**.
2. Select the tab for the resource assignment rule whose priority you want to change, and then click **Change Priority**.
3. In the box, change the priority.
4. Click **OK**.

Running a user resource assignment rule

For resources that have yet to be assigned to a threshold profile, run a user resource assignment rule to assign these resources to threshold profiles according to the assignment rule.

Procedure

1. In the **Administration** tab, select **Monitoring Settings > User Resource Assignment Rules**.
2. Select the tab for the resource for which you want to run the assignment rule, and then assign the target resource by clicking **Apply All Rules**.
3. Verify user resources to which threshold profiles have been set. Select **Monitoring Settings > User Resource Threshold Profiles**. Select the target threshold profile, and then click the **Target Resources** tab to display the resources list.

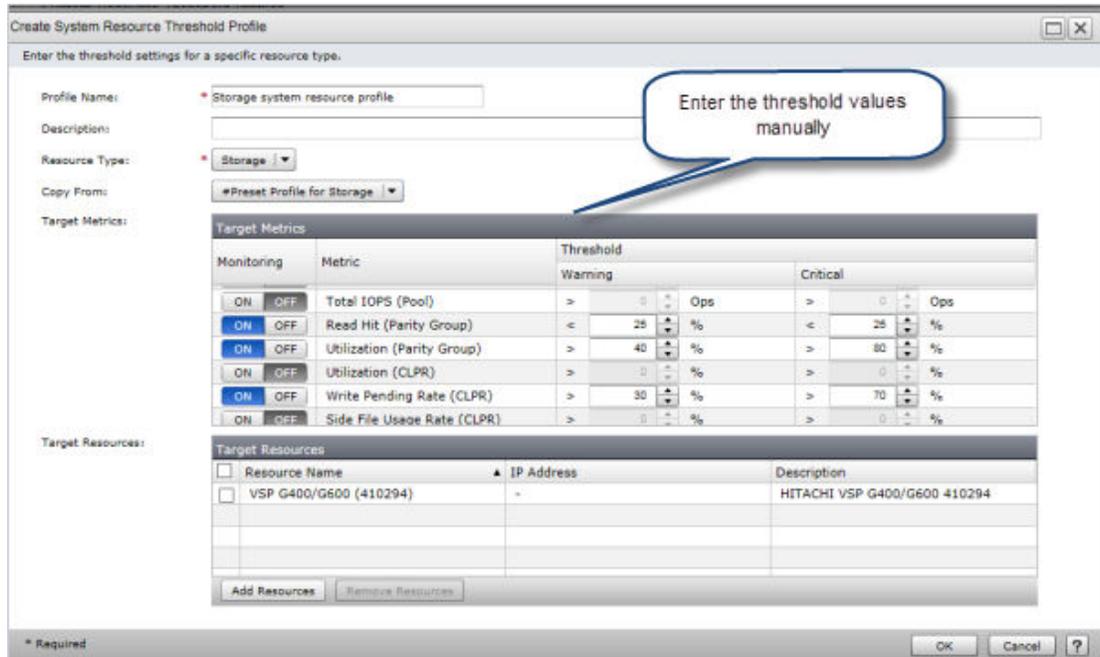
Setting thresholds for system resources

To detect hardware bottlenecks in system resources such as hypervisors and storage systems, create a profile and set thresholds.

Procedure

1. Go to the **Administration** tab, then from the navigation pane select **Monitoring settings > System Resource Threshold Profiles > Create Threshold Profile**.
2. In the **Create System Resource Threshold Profile** window, enter the profile name, description, and select the resource type. If required, copy the settings from the default profile or existing system resource profiles.

- Under threshold values, manually enter the threshold values for the performance metrics.



- Under **Target Resources**, click **Add Resources** to manually assign resources to the system resource threshold profile.

Removing monitored resources

If you remove resources from the environment or no longer require monitoring, you have the option of deleting the resources along with the collected performance data. In this case, all associated components displayed in the Ops Center Analyzer UI (such as CPU, memory, NIC, HBA, Disk, and VM for hypervisors) are removed, along with the performance data stored in the Analyzer detail view database.

When you remove the monitored resources, you must also delete the related probe instance. Otherwise, the resources are eventually detected and added back to the probe instance.

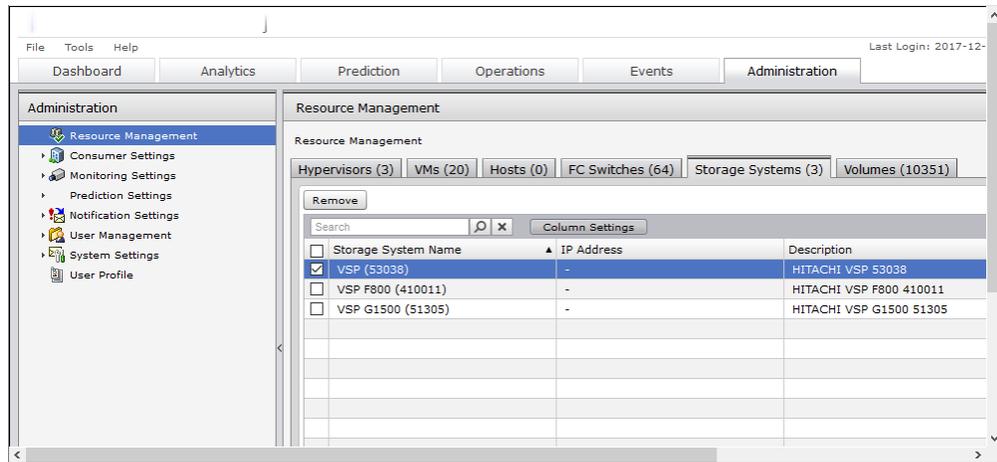
**Caution:**

- When a probe instance monitors multiple resources and you selectively remove a specific monitored resource, you must still delete the probe instance. When you delete the probe instance, the probe stops the data collection of all other associated resources. To resume the data collection of other associated resources, you must recreate the probe instance and add the resources.
- Ensure that you choose the correct resources to remove; the monitoring configuration and performance data is deleted and you cannot restore. If you remove the wrong resources by mistake and the probe instance is still in place, the resources are eventually detected and added again, but the historical performance data is lost, and you must redefine the consumer and threshold configuration. (Depending on the size of the configuration, the process of detecting and adding these resources may take a few hours.)
- Removing a probe instance and all the associated resources reduces the Analyzer detail view license node count.

To remove resources, you must first stop the probe, wait until the data transmission is complete, remove the monitored resources, and then delete the probe.

Procedure

1. Log on to the Analyzer probe server UI, and stop the probe from the **Status** window.
2. Wait until the data transmission is complete, which takes about 30 minutes. Verify the data import status of the probe instance that you want to delete on both Analyzer probe server and Analyzer detail view.
 - a. From the Analyzer probe server **Status** window, verify the last data collection time of **Performance Data** and **Configuration Data** for the probe.
 - b. From the Analyzer detail view dashboard, click the **Server Status** icon. From the **Status** window, verify the **Last Configuration Import Time** and **Last Performance Import Time** for the probe.
3. From the Ops Center Analyzer UI, click the **Administration** tab and select **Resource Management**.
4. Select the resource tab you want.
5. Select the check boxes for the resources you want to remove and click **Remove**.



6. The list of resources to be deleted appears in a new window. Click **OK**.
7. Because the removal cannot be undone, you are prompted to confirm your choice. Click **OK**.
8. To delete the probe, go to the Analyzer probe server UI, from the **Status** window, search for the probe that you want to delete, select the probe and then click **Delete**.

Result

The resources are removed in the background. If you have deleted a probe instance and the associated resources to reduce a license node count, you do not need to wait until the information has been updated to configure additional resources. You can later confirm that a license node has been released, as described in "Viewing and updating the Analyzer detail view license" in the *Hitachi Ops Center Analyzer Installation and Configuration Guide*.

Chapter 8: Managing consumers

Ops Center Analyzer simplifies management of infrastructure resources and the users of those resources. You can create consumers to associate users with the resources they are using.

Consumer settings

Create consumers to associate user resources with users of those resources.

Resources used by consumers, such as virtual machines or volumes, are grouped by company or business system consumer (also called management unit). Consumer definitions allow you to do the following:

- Search for consumer resources when the number of customer managed-resources increases.
- Verify which consumers are affected by failures on managed resources.
- Define the level of importance (grade) for each consumer to help you understand the problem severity and number of problems for each grade.

The following table is an example of creating four consumers based on the level of importance of business systems used by two companies.

In this example, four consumers are created.

Company name	Business system name	Importance of business system	Virtual machine name	Consumer name
Company A	Call center system	Very important	VM1	Company A (very important system)
			VM2	
			VM3	
	Email management system	Important	VM4	Company A (important system)
	Attendance management system	General	VM5	Company A (general system)

Company name	Business system name	Importance of business system	Virtual machine name	Consumer name
	Company SNS system	General	VM6	
Company B	Email management system	Important	VM7	Company B
	Sales management system	Important	VM8	

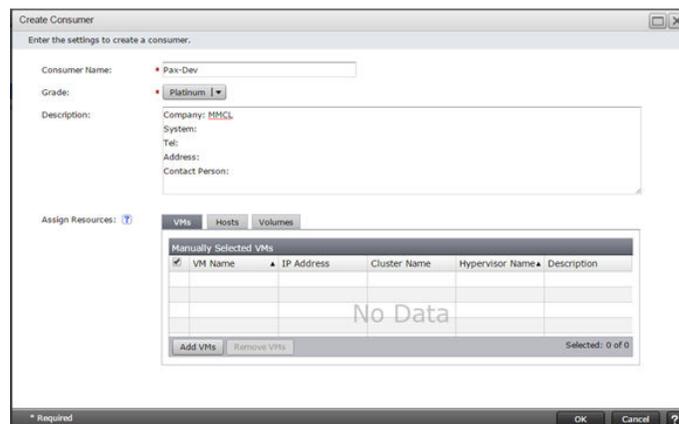
Creating a consumer

Create consumers to associate managed devices with the customers who use those devices.

Procedure

1. In the **Administration** tab, select **Consumer Settings > Consumers**.
2. Click **Create Consumer**.
3. Define the consumer name and grade, and then click **OK**.

The four default grades to assign consumers, in order of importance, are Platinum, Gold, Silver, and Bronze.



Creating a resource assignment rule

By creating resource assignment rules, the discovered resources are automatically registered to consumers 24 hours after discovery.

Procedure

1. In the **Administration** tab, select **Consumer Settings > Resource Assignment Rules**.

2. Select the tab for the resource for which you want to create an assignment rule, and then click **Create Rule**.
3. Define the assignment conditions and assignment destination.

Create VM Assignment Rule

Enter the rule name, and specify the required conditions. The rule is created with the lowest priority. To change the order of priority, go to the former screen.

Rule name: * MMCL

Description:

Resource Type: * VM (VMware) |

Condition: ?

Attribute	Condition	Value
VM Name	Equals	MMCL - +

Assign Destination: * Default Profile for VM (VMware) |

* Required

OK Cancel

4. Click **OK**.

Result

The created resource assignment rule appears in the resource assignment rules list.

Changing a resource assignment rule priority

Change the priority for a created resource assignment rule.

Procedure

1. In the **Administration** tab, select **Consumer Settings > Resource Assignment Rules**.
2. Select the tab for the resource assignment rule for which you want to change the priority, and then click **Change Priority**.
3. In the box, change the priority.
4. Click **OK**.

Running a resource assignment rule

Run a resource assignment rule to automatically assign resources to consumers.

Procedure

1. In the **Administration** tab, select **Consumer Settings > Resource Assignment Rules**.
2. Select the tab for the resource for which you want to run the assignment rule, and then assign the target resource by clicking **Assign VMs**, **Assign Hosts**, or **Assign Volumes**.
3. To verify that resources are assigned to consumers, select **Consumer Settings > Consumers**. Select a consumer, and then click the tab for the resource to display the resources list.

Creating multiple consumers (batch mode)

You can create multiple consumers in a single session using selectable defaults.

The Auto Create Consumers function generates a series of consumers associated with a resource type (VM, Host, or Volume). The consumer names are derived from a base parameter associated with the resource. For example, if you choose Volume as the resource type, the names can be based on a portion of the Volume Label, Host Group Name, or NVM Host Nickname. Furthermore, you can choose to export a consumer definition file so that you can make changes manually and upload the result to create your consumers. (See [About the consumer definition file format \(on page 113\)](#) for more information.)

Procedure

1. In the **Administration** tab, select **Consumer Settings > Consumers**.
2. To create consumers (or a new definition file), select **On demand**. If you already created a consumer definition file, select **Import consumer definition file**, click **Choose File** to select the file and skip to step 7.
3. Select a **Resource Type** to be associated with the consumers (VM, Host, or Volume). If you select "VM," you will see an additional option to **Assign the volumes associated with the VM to the same consumer** (default=off).
4. Select the "Base parameter" that will be used as a template for generating the consumer names.
5. For the "Name generation policy," choose the portion of the base parameter:
 - Use leading alphabetic portion
 - Use characters that precede separator
 - Use first N characters
 - Use entire base parameter

For example, if you choose **Host Group Name** for the base parameter, the consumer base name derived from the group name GAMMA47-HYDRA19 would be as follows:

- Use leading alphabetic portion: GAMMA
- Use characters that precede separator (default : "-"): GAMMA47
- Use first N characters (default: 8): GAMMA47-
- Use entire base parameter: GAMMA47-HYDRA19

Furthermore, all generated consumer names will include the prefix "_A_" (for example: _A_GAMMA).

6. Choose the default grade to be assigned to each consumer: Platinum, Gold, Silver, or Bronze.
7. Click **Preview** to see several examples of the consumer names that will be created.
8. If you are satisfied, click **Submit**. Otherwise, change the settings and preview again. You can also click **Export consumer definition file** to obtain the rule file created using your selections and make edits manually. You can then import the file as described in step 2.

9. Depending on the number of resources involved, the creation process takes time and proceeds in the background. A completion message is output to the event log when the process is complete. To examine the log, go to the **Events** tab and select the **All Events** tab. Choose **Filter** and search for "consumer."

Result

If the process fails, the log message will include any errors encountered along with suggested solutions.

About the consumer definition file format

The Auto Create Consumers function generates a definition file (CSV format) that can be edited and then imported to create consumers.

The following is a sample consumer definition file:

CONSUMER	GRADE	RESOURCE TYPE	ATTRIBUTE	CONDITION	VALUE	RULE NAME	RULE DESCRIPTION
_A_colossus-997	Platinum	Host	Host Name	Equals	colossus-997	_A_colossus-997_Host Name_Equals	_A_Host_Host Name_EntireBaseParam_Platinum_System_20220330100437
_A_forbin-761	Platinum	Host	Host Name	Equals	forbin-761	_A_forbin-761_Host Name_Equals	_A_Host_Host Name_EntireBaseParam_Platinum_System_20220330100437
_A_cpo-756	Platinum	Host	Host Name	Equals	cpo-756	_A_cpo-756_Host Name_Equals	_A_Host_Host Name_EntireBaseParam_Platinum_System_20220330100437

This example creates three consumers of Platinum grade based on hist names. In this case, the entire Host Name is used to name the consumer (Equals=Use entire base parameter).

The RULE NAME refers to the resource assignment rule that is automatically created for each consumer.

The RULE DESCRIPTION is simply the name given to the consumer definition file, in this case: `_A_Host_Host Name_EntireBaseParam_Platinum_System_20220330100437.csv`.

For best results, you should not make complex edits. The file is intended for making minor adjustments, such as:

- Setting a different grade for select consumers
- Making individual edits to consumer names
- Deleting unwanted consumers (by deleting an entire row)

If you need to make edits to the other columns, it is best to use the GUI.

Chapter 9: Setting notifications

Setting notifications is an easy way to stay informed on the status of infrastructure resources and events.

Monitoring resources is both an active and passive activity for IT administrators. Ops Center Analyzer allows you to configure email notifications that provide detailed information about problems with resource management. If multiple administrators rely on the Ops Center Analyzer service, you can create different profiles to deliver different types of information based on the profile settings.

Email notification settings

You can define who receives email notifications and when to send the notifications. Configure email notification settings to notify administrators about problems with management resources and provide details.

To enable email notifications, set up an email server, and then create condition profiles. You can edit and delete the created condition profiles.

Configuring the mail server

Configure the mail server and the email address of the sender to send emails in the following cases:

- To notify the administrator of problems that occur in monitored resources and information related to Analyzer server operations.
- To periodically send dashboard reports to users.

Before you begin

- Make sure you have Admin permissions for Ops Center Analyzer.
- Use the following settings for Email Notification and Send Test Mail:
 - Protocol: SMTPS, STARTTLS, cleartext
 - Authentication Methods: LOGIN, PLAIN, DIGEST-MD5

Procedure

1. In the **Administration** tab, select **Notification Settings > Email Server**.
2. Click **Edit Settings** to specify information about the mail server.
3. To verify that the mail server is configured correctly, click **Send Test Mail**.

4. Confirm that the test email arrives, and then click **Save Settings**.

Creating a condition profile

To receive emails from Ops Center Analyzer, create a condition profile and register email addresses.

Before you begin

You must have Admin permission of Ops Center Analyzer.

Procedure

1. In the **Administration** tab, select **Notification Settings > Notification Conditions**.
2. Select the **Condition Settings** tab, and then click **Create Notification Profile**.
3. Define the necessary items.
4. To add email addresses to which notifications are to be sent, click **Add Email Address** and then specify the email addresses. Select the check boxes for the email addresses for which you want to enable email notifications.

Change the status or descriptions of **Delivery Address** in the **Email Addresses** tab of the **Notification Conditions** window.

5. Click **OK**.

Enabling or disabling email addresses

Enable or disable registered email addresses to start or stop receiving email notifications.

Procedure

1. In the **Administration** tab, select **Notification Settings > Notification Conditions**.
2. Select the **Email Addresses** tab, and then select the check boxes for the email addresses to enable or disable the email notifications.
3. To enable email notifications, click **Activate**. To disable email notifications, click **Suspend**. Verify that the specified content is correct, and then click **OK**.

Chapter 10: Managing users

Add multiple user accounts to Ops Center Analyzer and set permissions based on predefined roles.

Ops Center Analyzer allows you to create multiple user accounts with different permissions. For security, enable user account locking and password protection.

User management

Create a user account in Ops Center Analyzer, and then set permissions by the user role. When you configure Ops Center Common Services for Ops Center Analyzer, you can manage users from the Ops Center portal. For more information about how to configure Common Services for single sign-on, see the Hitachi Ops Center Analyzer Installation and Configuration Guide.

User account types

After installation, a default system account is set up. The system account (user name: system; default password: manager) is a fully privileged, built-in administrator account used to manage all the functions in Ops Center Analyzer. Use the system account to do all tasks and manage users in Ops Center Analyzer. You cannot delete or change the default user ID.

When registering new user accounts, set permissions for the types of tasks each user can do based on the user role.

After adding basic user information such as username, password, email, and description, set permissions for available applications, such as:

- User Management
- IAA

User permissions

There are four types of user permissions:

- Admin
Users with Administrator permissions can perform all management tasks in Ops Center Analyzer except those related to managing users, configuring security options, and modifying I/O control settings.
- Modify
Users with Modify permissions can perform all management tasks in Ops Center Analyzer except those related to managing users, configuring security options, setting up email notifications, configuring the connection settings for the Analyzer detail view server, and modifying I/O control settings.
- User Management
Users with User Management permissions can perform all tasks related to managing users and configuring security options. You can use this type of permission with other types of permissions.
- StorageOps
Users with StorageOps permissions can perform all tasks related to the I/O control settings on the Operations tab. These users can monitor and regulate the I/O rate of the monitored volumes. Only users with Admin or Modify permissions can be assigned StorageOps permissions.

User account administration

You can do the following tasks on the **User Management** window. Select User Management > Users and Permissions.

- Add or delete user accounts.
- Set permissions for user accounts.
- Reference or edit user account profiles.
- Change the password of a user account.
- Change the status of a locked user account.



Note: Before deleting the users from the User Management window, make sure you delete the custom reports created by those users.

User account authentication

If you have configured Ops Center common services for using single sign-on, you can manage Ops Center Analyzer user accounts from the Ops Center portal.

You can also manage user accounts by linking to an external authentication server, such as an Active Directory server, LDAP directory server, RADIUS server, or Kerberos server.

However, the built-in accounts (system) cannot be authenticated on an external authentication server. The Ops Center Analyzer user account used to connect to external authentication/authorization servers is managed as an Active Directory (Authorization) group. Permissions that are specified for authorized groups are also applied to users who belong to nested groups.

User permissions for Hitachi Ops Center Automator services

If you have Ops Center Automator in your infrastructure environment, you can assign access privileges to Ops Center Analyzer users.

To run action and service templates, you need to set up permission in the common component for each product. Ops Center Analyzer users need to be in the Ops Center Automator User Group and Service Group. Likewise, Ops Center Automator users need modify or admin permissions in Ops Center Analyzer.



Note: While you can assign the Ops Center Analyzer account users to any user group in Ops Center Automator, the best practice is to assign Ops Center Analyzer account users to the **SubmitGroup**.

Configuring external authentication for users

You can set up external authentication systems for user logins.

Before you begin

- The Ops Center Analyzer server must be linked to an external authentication server.
- The Ops Center Analyzer server must be configured to support user authentication, which activates the Change Auth button in the GUI, and which presents authentication options such as Internal for a local account, or LDAP for external authentication.
- The Ops Center Analyzer user ID must exist on the external authentication server. It is recommended that user ID information be acquired from the external authentication server administrator before creating accounts.

Procedure

1. From the **Administration** tab, select **Users and Permissions**.
2. Select the **Users** folder, then select one or more users (using the check box) whose authentication method you want to change, or click **Add User** to create a new account.

**Note:**

When creating a new account, only the **User ID** is required for external authentication, and must match a user ID on the external authentication server. For a local (internal) account, a **User ID** and **Password** are both required. When external authentication is available, new user accounts created without a password value are automatically configured to use external authentication (for example, LDAP is selected for you). Fill in the desired fields, and click **OK** to create the user account.

If you are completing an external authentication in a multiple-domain configuration or by using realms, specify a user ID that includes the domain name or realm name for the **User ID**. Example: *user-name@domain-name* or *realm-name*.

When you log in, also specify your user ID in this format.

3. If you have selected existing users, click **Change Auth**. From the list, select the authentication method (for example, LDAP) and click **OK**. The user list is redisplayed.
4. Review the **Authentication** column to verify the authentication method.

Result

On the next login attempt by each user, the users' login credentials (user ID and password) will be validated using the external authentication server.

Configuring external authentication for groups

Set up external authentication systems for user groups.

When linking with an external authentication server, if using together with Active Directory as an external authorization server, user permissions can be managed by using the Active Directory groups (authorization groups) registered on the external authorization server. In this case, user permissions are specified for each group.

Before you begin

- The server must be linked to an external authentication (authorization) server.
- The Ops Center Analyzer server must be configured to support group authentication, which activates the Groups folder in the GUI.
- The Ops Center Analyzer user group must exist on the external authentication (authorization) server. It is recommended that domain and group information be acquired from the external authentication server administrator.

Procedure

1. From the **Administration** tab, select **Users and Permissions**.
2. Click the **Groups** folder to display the **Domain List**. This is a list of external authentication servers listed by domain name, and host name or IP address. If the **Groups** folder is not displayed, see the prerequisites above.

3. Select the desired **Domain Name** to display the **Group List**, which may be empty ('No Groups' is displayed). Click **Add Groups**.
4. Enter the **Distinguished Name** for the group. Use **Check DN** to verify a correct DN entry. Click **Ok** to save your group and redisplay the **Group List**. Note that the **Group Name** is derived from the entered DN. To specify multiple groups, note that:
 - You can add multiple DNs at the same time using the "+" button.
 - If multiple DNs are listed, you can remove an entry with the "-" button.
 - **Reset** clears all DN entries.
5. From the **Group List**, click the **Group Name** link, then click **Change Permission** and set the Ops Center Analyzer permissions for the group (repeat this step for each new group).
6. Your groups will now be visible from the **Administration** tab, **User Groups**. You can affiliate the groups with resource groups and roles, just like Ops Center Analyzer user groups. If you delete external authentication groups from **Users and Permissions** at a later time, the groups are also removed from the **User Groups** list.

Result

On the next login attempt by each group member, the login credentials (user ID and password) will be validated using the external authentication (authorization) server.



Tip:

To delete registered authorization groups, select the check boxes of the groups to be deleted, and then click **Delete Groups**.

Security

Set the appropriate security options for user logon.

You can do the following operations in Ops Center Analyzer :

- Set a password policy

To prevent third parties from guessing passwords, define conditions such as a minimum number of characters required for a password and the type of characters to include in a password.
- Configure settings to automatically lock accounts

If repeated attempts are made to log on to a user account using an incorrect password, the user account is locked to prevent unauthorized access.
- Set warning banners

As a security measure for logon, you can customize the message (warning banner) to display in the logon window.

When you are in the Administration tab, you can do these operations in the window that opens when you select User Management > Security.

Chapter 11: Additional dashboard reports

Use additional dashboard reports to collect information for statistical analysis and performance trends.

Ops Center Analyzer offers numerous types of reports that provide information on various resources.

System Status Summary for Storage Resources

The System Status Summary for Storage Resources report displays the performance status summary of your monitored volumes and system resources, such as SAN and storage components.



System Status Summary for Volumes

The **Volume** pane displays an information gauge chart using three colors to depict the severity levels of alerts. Green indicates normal status, red indicates critical status, and yellow indicates warning status. The chart also displays the number of volumes that triggered alerts out of the total number of monitored volumes.

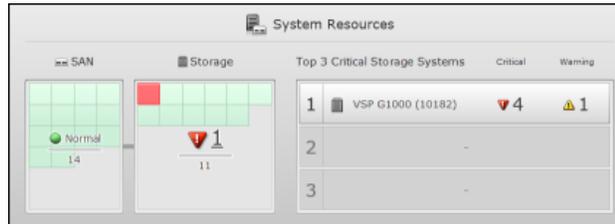


Analyze the status of monitored volumes based on the key performance metrics such as Response Time, IO Rate, and IOPS. The bar graph displays the number of volumes that triggered critical or warning alerts when the value of any monitored metric exceeds the defined threshold values.

To view details about the volumes that exceeded the defined critical or warning thresholds, click the number link in the information gauge chart or bar graph. A list of affected volumes appears in a new window. Select the volume and then click Show E2E View to view the data center topology and review the system configuration.

System Status Summary for System Resources

The System Resources pane displays the following reports:



A heat map uses three colors to depict the severity levels of alerts triggered by the SAN (switches) and storage components (ports, processors, cache, pools, and parity groups). The green tiles indicate normal status, red tiles indicate critical status, and yellow tiles indicate warning status. The heat map also displays the number of SAN and storage components associated with alerts. To view details about the resources that exceeded the defined critical or warning thresholds, click the number link in the heat map. A list of affected SAN and Storage resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

A list displays the top three critical storage systems with critical and warning alerts. To view details about the storage system components that exceeded the defined thresholds, click the number link. A list of components associated with the selected storage system appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

System Status Summary for Infrastructure Resources

The System Status Summary for Infrastructure Resources report displays the performance status summary of monitored user resources (such as VMs, hosts, and volumes) and system resources (such as server, SAN, and storage components).



System Status Summary for User Resources

The System Status Summary - VM/Host and Volume reports display the status of all monitored VMs, hosts, and volumes.



Both the VM/Host pane and the Volume pane display a Resource Status information gauge, where the top number is the total critical or warning alerts received from the VMs and hosts, or volumes that exceeded the critical or warning thresholds for any monitored metric. The bottom number indicates the total number of VMs, hosts, or volumes in the system.

Under Metrics, a bar graph displays the total number of VMs and hosts with critical and warning alerts for any monitored metric.

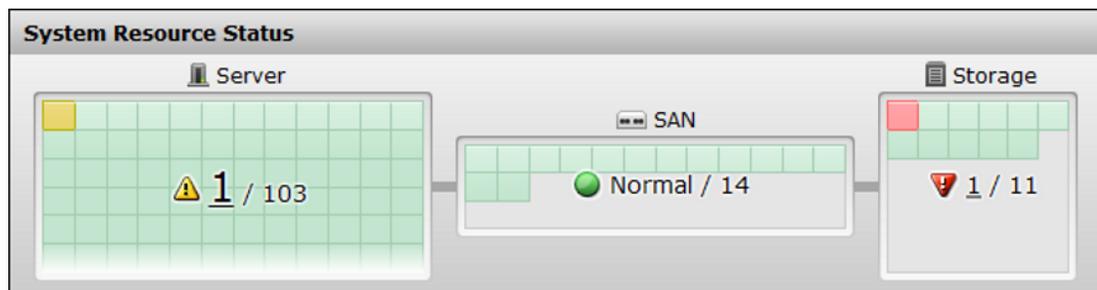
For the Volume summary, Metrics displays a bar graph of the total number of volumes with critical and warning alerts for any metrics. In the bar graph, red is the total number of critical alerts and yellow is the number of warning alerts.

For example, if there are 3 critical alerts and 5 warning alerts, then both critical and warning display in the gauge. If the number of critical alerts is greater than 9, only the red bar (critical) displays because the maximum value of the gauge is 10.

To view details about the resources that exceed the defined critical or warning thresholds, click the number link in the information gauge chart or bar graph. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration to analyze the performance problem.

System Resource Status report

The System Resource Status report is one of the default reports that appears on the Ops Center Analyzer dashboard. It provides a heat map of the current status of system resources such as server (CPU, memory, NIC, HBA, and disk), SAN (switches), and storage (ports, processors, cache, pools, and parity groups) components.



Each red tile shows a critical alert, and each yellow tile shows a warning alert. Unknown resources are considered Normal and are represented by green tiles.

To view details about the resources that exceeded the defined thresholds, click the number link. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Consumer reports

Consumer reports displays the performance status of the consumers based on the performance of the VMs and hosts as well as volumes assigned to them.

Consumer Summary report

The Consumer Summary report on the Ops Center Analyzer dashboard provides a summary of the status of consumers. It displays the total number of consumers that are affected by critical and warning alerts from VMs, hosts, or volumes assigned to consumers. This report also displays a table by consumer name, grade, and status for each alert.

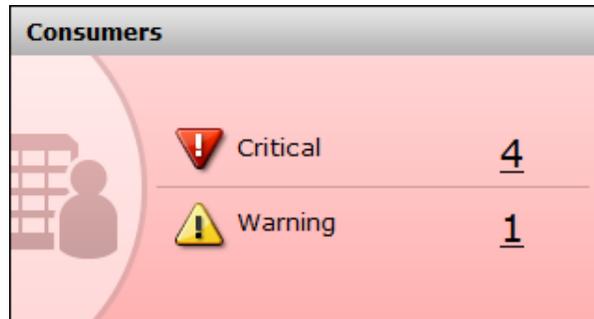
To view details about the affected consumers, click the number link in the left pane.

Consumer Summary			Consumer Name	Grade	Status
  4  1	Company A import...	Platinum	Critical		
	Company B	Gold	Critical		
	Company Y	Bronze	Critical		
	Company Z	Bronze	Critical		
	Company C	Gold	Warning		

The Consumers window lists the consumer name, consumer grade, overall status, and the detailed status of each assigned user resource. To view details about the resources that exceeded the defined thresholds, click the number link. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Consumers report

The Consumers report on the Ops Center Analyzer dashboard displays the total number of consumers that are affected when critical and warning thresholds are exceeded for VMs, hosts, or volumes assigned to consumers. To view details about the affected consumers, click the number link in the pane that shows the critical and warning alerts. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.



Consumers		
	Critical	<u>4</u>
	Warning	<u>1</u>

The Consumers window lists the following:

- Grade: The grade associated with the consumer.
- Status: The most severe status for all user resources assigned to the consumer.
- Summary status: The most severe status for VMs, hosts, and volumes of the affected consumer.

Consumers - Critical report

The **Consumers - Critical** report on the Ops Center Analyzer dashboard provides the total number of consumers that are affected by critical thresholds exceeded by VMs, volumes, or hosts assigned to consumers.



Consumers - Critical	
	<u>4</u>

To view details about the affected consumers, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Consumers - Critical Each Grade report

The **Consumers - Critical Each Grade** report on the Ops Center Analyzer dashboard displays the number of platinum, gold, silver, and bronze, or others that are affected by critical alerts in any of their managed resources, such as volumes, VMs, and hosts. This report is useful if you want to know which consumers are affected by critical alerts based on grades.

Consumers - Critical Each Grade	
☆☆☆ Platinum	<u>1</u>
☆☆ Gold	<u>1</u>
☆ Silver	0
Others	<u>2</u>

To view more details about the affected consumers, click the number link next to the grade. The **Consumer** window lists the consumer name, consumer grade, overall status, and detailed status of each assigned user resource.

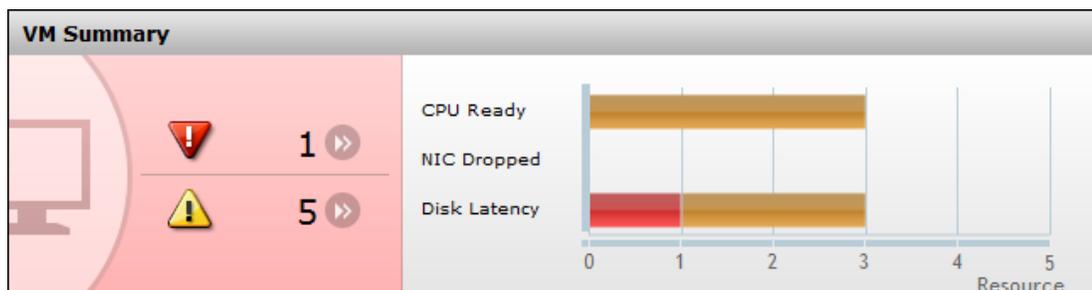
To view details about the resources that exceeded the defined critical thresholds, click the E2E View on the **Consumers** window. You can view the data center topology, and review the configuration and status information.

VM reports

The VM reports display the performance of the monitored VMs based on the key performance indicators such as, CPU ready, disk latency, and NIC dropped.

VM Summary report

The VM Summary report on the Ops Center Analyzer dashboard provides a summary of the performance and status of your registered virtual machines.



If the numbers are above 0, VMs exceeded the critical and warning thresholds. The horizontal bar graph provides a breakdown of the overall VM status, which indicates the total number of critical and warning alerts compared to the number of affected resources from the following monitored metrics:

- CPU ready: The CPU metric SLO status of all your monitored VMs.
- NIC dropped: The NIC metric which measures the rate of IP packets drops sent or received by the NIC.
- Disk latency: The read and write latency status of all your monitored VMs.

Hover over the bar graph to view the VM alert status. The status indicates how many resources are affected by critical and warning alerts.

To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Data is not collected from VMs on a hypervisor for the following metrics:

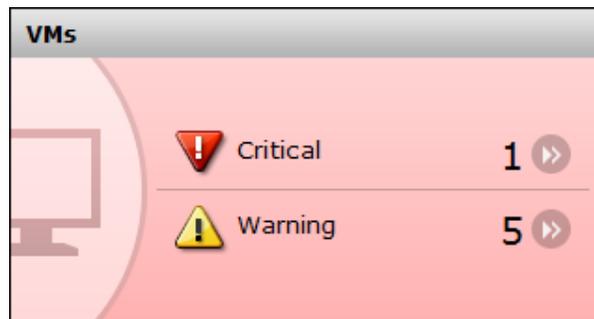
- CPU Ready
- NIC Dropped
- Disk Latency

As a result, data about VMs is not included in the aggregate results for the following Dashboard reports:

- System Status Summary
- VM Summary
- VM CPU Ready
- VM NIC Dropped
- VM Disk Latency

VMs report

The VMs report on the Ops Center Analyzer dashboard displays the total number of VMs with critical and warning alerts.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

VMs - Critical

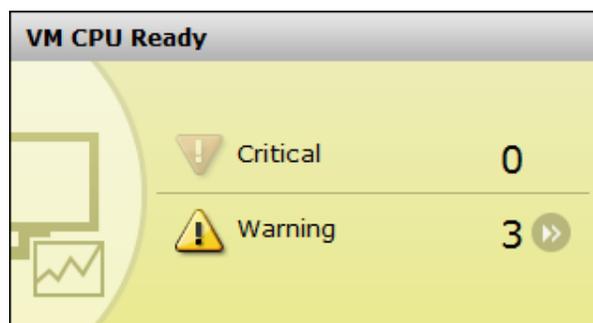
The VMs - Critical report on the Ops Center Analyzer dashboard displays the number of virtual machines that exceeded the defined critical threshold for any monitored metric.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

VM CPU Ready report

The VM CPU Ready report on the Ops Center Analyzer dashboard displays the CPU metric SLO status of all your monitored virtual machines. The CPU metric refers to the ratio of time that the virtual machine was ready but could not be scheduled to run on a physical CPU. The numbers indicate the total critical and warning alerts received based on the relevant thresholds set for the CPU metric.

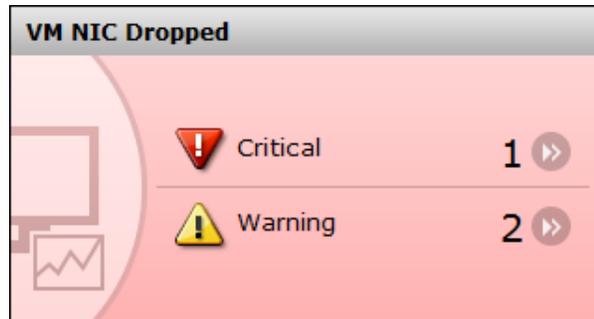


All critical and warning alerts directly affect the SLO, so start analyzing the problem quickly.

To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

VM NIC Dropped report

The VM NIC Dropped report on the Ops Center Analyzer dashboard displays the IP packets sent or received IP packets that were dropped for all your monitored virtual machines. The numbers indicate the total critical and warning alerts received compared to the relevant thresholds set for the NIC metric. The metric refers to sent or received packets that were dropped of the virtual NIC.



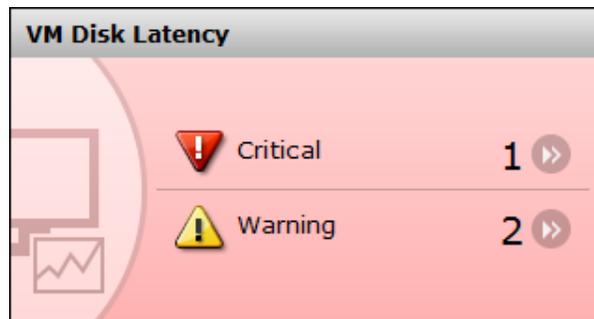
All critical or warning alerts directly affect the SLO, so start analyzing the problem quickly.

To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

VM Disk Latency report

The **VM Disk Latency** report on the Ops Center Analyzer dashboard displays the disk read and write latency status of all your monitored virtual machines. The numbers indicate the total critical and warning alerts received compared to the relevant thresholds set for the disk latency metric. The disk latency metric refers to the read and write latency to and from the virtual machine disk.

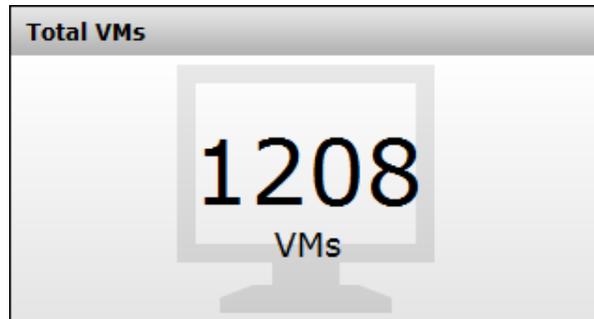
All critical or warning alerts directly affect the SLO, so start analyzing the problem quickly.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Total VMs report

The Total VMs report on the Ops Center Analyzer dashboard displays the total number of monitored virtual machines.



If you plan to register VMware vCenter server, then all VMs connected to this server are monitored.

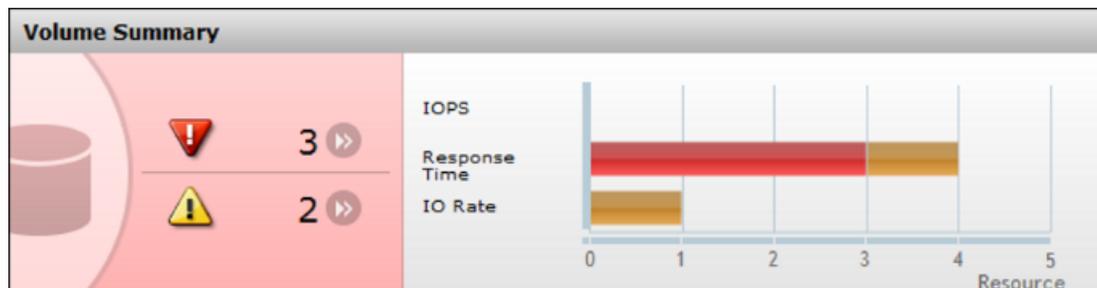
Volume reports

The volume reports displays the performance of your storage system volumes based on the key performance indicators such as IOPS, response times, and I/O rate.

Volume Summary report

The **Volume Summary** report on the Ops Center Analyzer dashboard provides an overview of the performance of your storage system volumes. It shows the total number of volumes with critical and warning alerts and also a summary of alerts received for the following three monitored metrics:

- IOPS: I/O operations per second of the monitored volume.
- Response time: Average response time of the volume based on the I/O count.
- IO rate: I/O operations transfer rate.

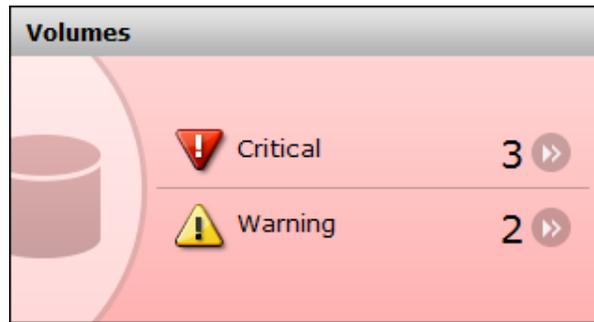


Hover over the bar graph to view the alerts type and number of volumes with critical and warning alerts for these three metrics.

To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Volumes report

The **Volumes** report on the Ops Center Analyzer dashboard displays the total number of volumes with critical and warning alerts.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Volumes - Critical report

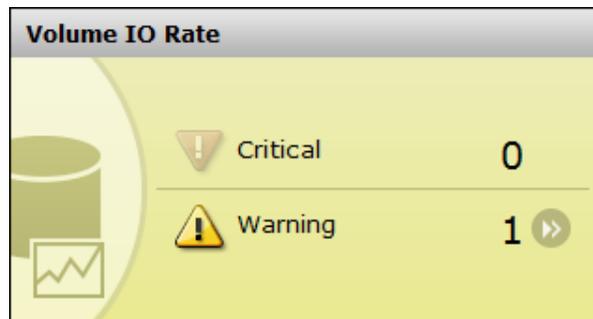
The **Volumes - Critical** report on the Ops Center Analyzer dashboard displays the total number of volumes that exceeded the critical threshold for any metric.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Volume IO Rate report

The **Volume IO** rate report on the Ops Center Analyzer dashboard displays the total number of critical and warning alerts received from volumes that exceeded the critical and warning thresholds for the I/O rate metric.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Volume IOPS report

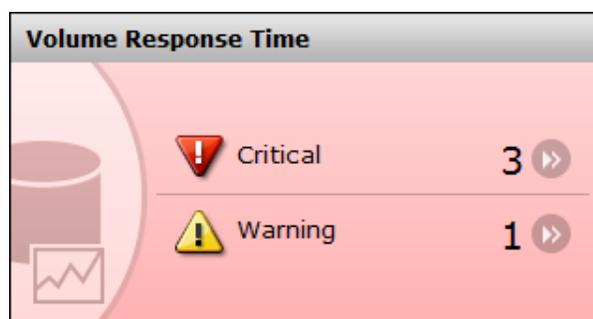
The **Volume IOPS** report on the Ops Center Analyzer dashboard displays the total number of critical and warning alerts received from volumes that exceeded the critical and warning thresholds for the IOPS metric.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Volume Response Time report

The **Volume Response Time** report on the Ops Center Analyzer dashboard displays the response time status for all monitored volumes. The numbers indicate the total number of critical and warning alerts received from volumes that exceeded the critical and warning thresholds for the response time metric.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Total Volumes report

The **Total Volumes** report on the Ops Center Analyzer dashboard displays the total number of volumes.

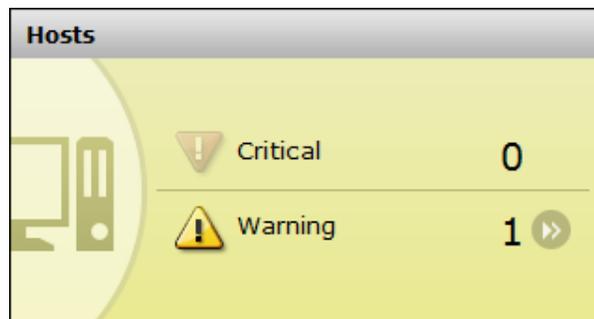


Host reports

The host reports display the total number of hosts monitored and the status of the monitored hosts.

Hosts report

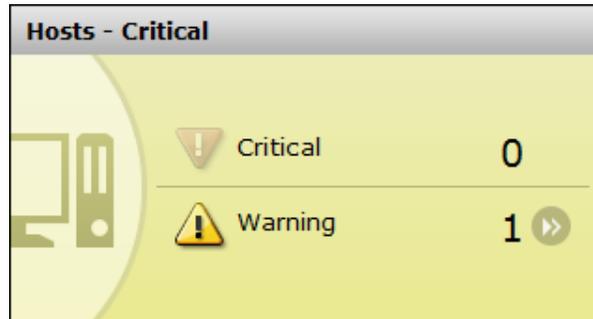
The **Hosts** report on the Ops Center Analyzer dashboard displays the total number of hosts that have critical and warning alerts.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Hosts - Critical report

The **Hosts - Critical** report on the Ops Center Analyzer dashboard displays the number of hosts that exceeded the defined critical threshold for any monitored metric. The host is monitored with the thresholds defined with Windows-based profiles in the user resource thresholds.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Total Hosts report

The **Total Hosts** report on the Ops Center Analyzer dashboard displays the total number of hosts monitored in the environment.



User resource reports

The user resource reports display the total number of monitored user resources such as volumes, VMs, and hosts.

Total User Resources report

The **Total User Resources** report on the Ops Center Analyzer dashboard displays the total number of monitored user resources. It also shows a breakdown of the total number for each resource type:

- Volumes
- VMs
- Hosts

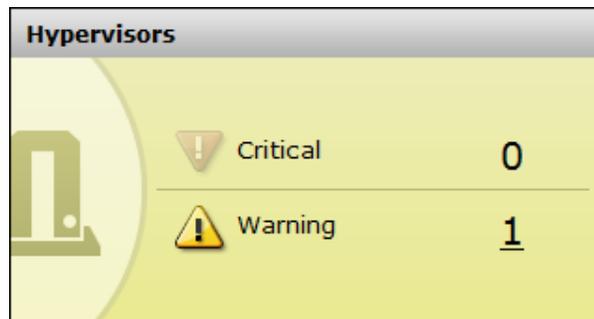


Hypervisor reports

The hypervisor reports display the total number of monitored hypervisors and the status of the monitored hypervisors.

Hypervisors report

The **Hypervisors** report on the Ops Center Analyzer dashboard displays the total number of monitored hypervisors that have critical and warning alerts.



To view details about the affected hypervisors, click the number link in the critical and warning alerts pane to display the following on the Hypervisors window:

- Hypervisor Name: The name of the hypervisor.
- IP Address: The IP address of the hypervisor.
- Status: The most severe status for hypervisor.
- Cluster Name: The cluster name of the hypervisor.
- Description: A short description of the hypervisor.

To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Total Hypervisors report

The **Total Hypervisors** report on the Ops Center Analyzer dashboard displays the total number of monitored hypervisors.

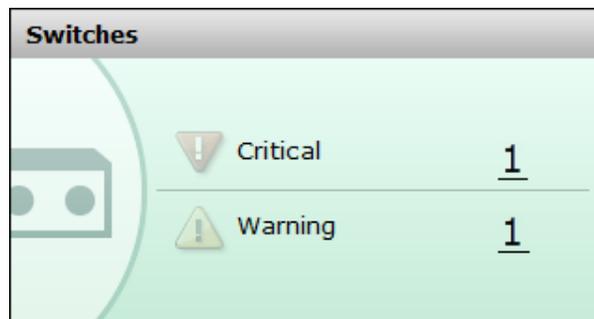


Switch reports

Switch reports display the total number of monitored switches and status of monitored switches.

Switches report

The **Switches** report on the Ops Center Analyzer dashboard displays the total number of FC switches that have critical and warning alerts.



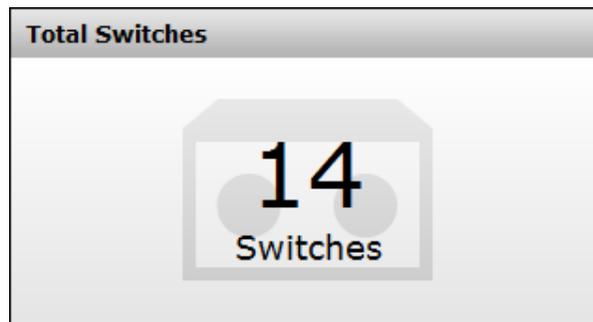
To view details about the affected switches, click the number link in the critical and warning alerts pane to display the following on the **Switches** window:

- Switch Name: The name of the switch.
- IP Address: The IP address of the switch.
- Status: The most severe status for the switch.
- Description: A short description of the switch.

To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Total Switches report

The **Total Switches** report on the Ops Center Analyzer dashboard displays the total number of monitored FC switches.



Storage system reports

Storage Systems report

The **Storage Systems** report on the Ops Center Analyzer dashboard displays the total number of storage systems that have critical and warning alerts.



To view details about the affected storage systems, click the number link in the critical and warning alerts pane to display the following on the **Storage Systems** window:

- Storage System Name: The name of the storage system.
- IP Address: The IP address of the storage system.
- Status: The most severe status for the storage system resource.
- Description: A short description of the storage system.

To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Total Storage Systems report

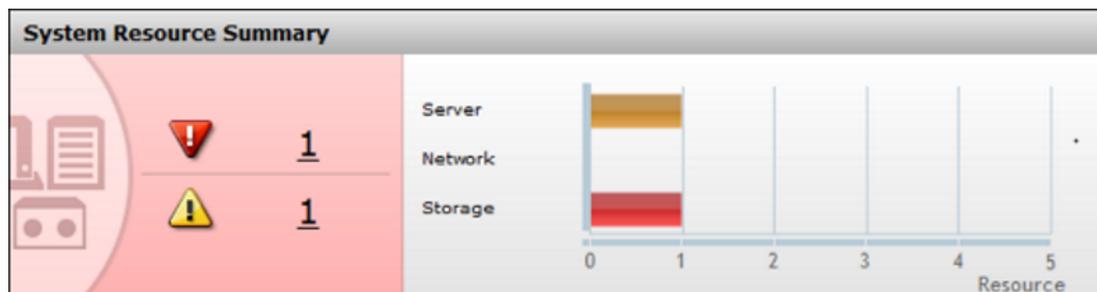
The **Total Storage Systems** report on the Ops Center Analyzer dashboard displays the total number of monitored storage systems.



System resource reports

System Resource Summary report

The **System Resource Summary** report on the Ops Center Analyzer dashboard provides a summary of the performance and status of your registered system resources such as hypervisors, switches, and storage systems. The report provides the total number of critical and warning alerts, and also includes these numbers in a bar graph.



If the numbers are higher than 0, the system resources exceeded the critical and warning thresholds. The bar graph provides a summary of the total number of resources with critical and warning alerts, by each resource type:

- Server
- Network
- Storage

Hover over the bar graph to view the system resource alert status. The status indicates how many system resources are affected by critical or warning alerts. To view details about the resources that exceeded the defined critical and warning thresholds, click the number link.

The System Resource window lists the resource name, IP address, status, device type, and description of each assigned system resource. On the System Resource window, click E2E View to view the data center topology and review the configuration and status information.

System Resources - Critical report

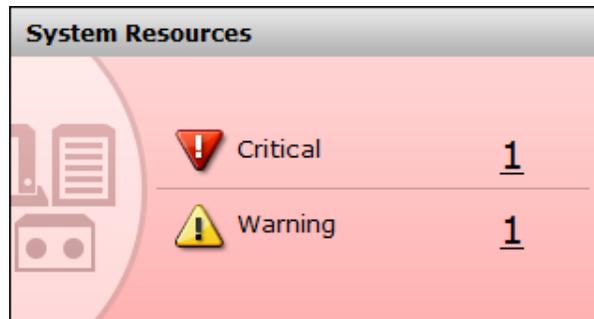
The **System Resources - Critical** report on the Ops Center Analyzer dashboard displays the number of system resources that exceeded the defined critical threshold for any monitored metric.



To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

System Resources report

The **System Resources** report displays the total number of monitored system resources, such as hypervisors, switches, and storage systems, that have critical and warning alerts.



To view details about the affected system resources, click the number link in the pane that shows the critical and warning alerts to open the **System Resources** window with the following:

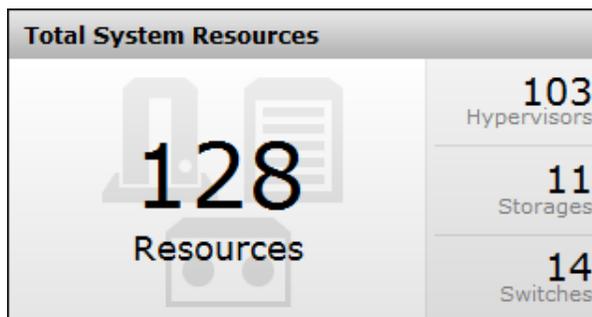
- Resource Name: The name of the system resource.
- IP Address: The IP address of the system resource.
- Status: The most severe status of the system resource.
- Description: A short description of the system resource.

To view details about the affected resources, click the number link in the report. A list of affected resources appears in a new window. Select the resource and then click Show E2E View to view the data center topology and review the system configuration.

Total System Resources report

The **Total System Resources** report on the Ops Center Analyzer dashboard displays the total number of monitored system resources. It also shows a breakdown of the total number of each resource type:

- Hypervisors
- Storage systems
- Switches



Event reports

System Event Status report

The **System Event Status** report on the Ops Center Analyzer dashboard displays the number of critical, warning, and information events for Management or Event Action events that occurred in the past 24 hours. An example of a Management event is a service start event. An Event Action event is a return code notification of the executed action. For example, when launching an event action the command fails, a failure notification is sent. These types of events are generated when the system settings need to be verified or configured.



To view more details for each event, see the Events > System Events tab.

Resource Event Status report

The **Resource Event Status** report on the Ops Center Analyzer dashboard displays the number of critical, warning, and informational events for resource events that were generated in the past 24 hours. The events are generated when a device or component such as servers or network devices has a problem or does not perform optimally.

Resource Event Status		
	Critical	4
	Warning	7
	Information	25

To view more details for each error, see the Events > Resource Events tab.

Capacity reports

Ops Center Analyzer enables consolidated storage capacity reporting by consumers, drive types, and pools across all storage systems. By analyzing the storage capacity reports you gain insight into your storage capacity, identify unused capacity, and track storage usage. Monitor the capacity reports to forecast the capacity requirements for your data center and plan new purchases of storage systems, if necessary.

The capacity reports display the latest capacity information in the past 24 hours. You can also view the total capacity of multiple resources and all related resources in the past 24 hours. For example, the report Recent Capacity by Consumer aggregates not only volumes that are currently associated with a consumer, but also volumes that are removed from the consumer within the past 24 hours.

Capacity of VMware Datastores report

The **Capacity of VMware Datastores** report displays a list of all monitored VMware datastores along with capacity and usage.

Capacity of VMware Datastores			
Name	Capacity (GB)	Used Capacity	FreeSpace (GB)
JH-23519_R600#1	99.75	75.97 GB	23.78
JH-23519_R700#2	99.75	75.97 GB	23.78
JH-23519_HM700#2	99.75	83.55 GB	16.20
datastore_DF850	99.75	0.95 GB	98.80
JH-28587_R800_10	99.75	0.95 GB	98.80

Capacity of VMware Datastores with Usage Exceeding 80% report

The **Capacity of VMware Datastores with Usage Exceeding 80%** report displays a list of datastores with high usage so that you can monitor the remaining capacity and allocate additional storage.

Capacity of VMware Datastores with Usage Exceeding 80%					
Name ▲	Type	Capacity	Used Space	Space Utilizatio..▼	Free Space
JH-23519_HM...	VMFS	99.75 GB	83.55 GB	83.76 %	16.20 GB
bs2k8	VMFS	199.75 GB	185.16 GB	92.70 %	14.59 GB
JH-27635_HM...	VMFS	199.75 GB	188.92 GB	94.58 %	10.83 GB
bs2k7	VMFS	199.75 GB	169.81 GB	85.01 %	29.94 GB

Capacity by Storage System report

The **Capacity by Storage System** report provides the breakdown of the overall storage capacity indicating capacity from physical parity groups and DP pools. Using this report you can monitor the capacity consumption by storage systems and optimize the use of storage resources. You can track the capacity savings achieved by using advanced data reduction technologies such as deduplication and compression, and FMD (Flash Module Drive) enabled for accelerated compression.

Capacity by Storage System								
Model ▲	Serial No ▲	Thin Total	Thin Used	Thin Free	Subscribed Total	FMC Capacity Saving	DKC Comp/Dedupe Ca...	
VSP F800	418923	64053.98 GB	2556.52 GB	61497.49 GB	760642.89 GB	0.00 GB	498.00 GB	

The report contains the following metrics:

- Model: The storage system type.
- Serial No: The serial number of the storage system.
- Thin Total: The sum of the capacity of DP pools in a storage system.
- Thin Used: The sum of the used capacity of DP pools in a storage system.
- Thin Free: The sum of the free capacity of DP pools in a storage system.
- Subscribed Total: The sum of the capacity of DP volumes in a storage system.
- FMC Capacity Saving: The total capacity savings of DP pools achieved by using FMD that supports accelerated compression in a pool.
- DKC Comp/Dedupe Capacity Saving: The total capacity savings of DP pools achieved by using the DKC capacity saving feature in the storage system.
- Total Parity Group Free: The sum of the free capacity of parity groups in a storage system.

**Note:**

The following attributes are not displayed for concatenated parity groups. From the Details windows, check the attributes of the individual parity groups that make up a concatenated parity group.

- Capacity
- Drive Status
- Number of used spare drives
- SSD Used Endurance Indicator
- FMD Battery Life Indicator

Click the link to the storage system serial number to open the Analyzer detail view UI in a separate browser. You can analyze the storage system performance reports on the Reports pane.

Capacity by Storage System report (HUS 100 Storage Systems)

The **Capacity by Storage System** report for Hitachi Unified Storage 100 (HUS 100) series provides the breakdown of storage capacity indicating capacity from physical parity groups and DP pools. Use this report to monitor the storage capacity consumption and optimize the use of storage resources.

Model	Serial No	Thin Total	Thin Used	Thin Free	Subscribed Total
HUS130	921103110	7280.00 GB	4651.00 GB	2629.00 GB	6142.40 GB

The report displays the following metrics:

- Model: The storage system type.
- Serial No: The serial number of the storage system
- Thin Total: The sum of the capacity of DP pools in a storage system.
- Thin Used: The sum of the used capacity of DP pools in a storage system.
- Thin Free: The sum of the free capacity of DP pools in a storage system.
- Subscribed Total: The sum of the capacity of DP volumes in a storage system.
- Total Parity Group Free: The sum of the free capacity of parity groups in a storage system.

Click the link to the storage system serial number to open the Analyzer detail view UI in a separate browser. You can analyze the storage system performance reports on the Reports pane.

Saving Ratio by Storage System report

The **Saving Ratio by Storage System** report provides a breakdown of the capacity saved through data reduction functions such as controller-based compression, deduplication, and accelerated compression.

Using this report, you can monitor the amount and ratio of capacity saving in a pool or per storage system.

This report applies to VSP 5000 series, VSP E series, VSP G1x00, VSP Gx00 models, VSP F series, and VSP N series .

The report contains the following metrics:

- Model: The storage system type.
- Serial No: The serial number of the storage system.
- Saving Capacity: The amount of saved logical capacity against the used physical capacity.
- Saving Ratio: The ratio of used logical capacity to used physical capacity.

If any of the data reduction functions are enabled, the pool capacity categories (total capacity, used capacity, and free capacity) displayed in the capacity reports indicate the values after capacity saving. If the pool contains a parity group for which accelerated compression is enabled, then the values before compression are displayed.

To check the capacity of a pool containing an FMC drive, you typically also need to check the physical capacity. However, if the pool is set up to automatically add pool volumes in these storage systems, you only need to check the physical capacity.

The supported storage systems also allow the addition of pool volumes automatically (according to the compression ratio of the parity group) when **Autoexpansion** is enabled.

Total Efficiency by Storage System report

The **Total Efficiency by Storage System** report provides a summary of the capacity saved without system data through data reduction functions such as controller-based compression, deduplication, accelerated compression, snapshot efficiency, and provisioning efficiency.

Using this report, you can monitor the amount and ratio of capacity saving per storage system.

This report applies to VSP 5000 series, VSP E series, VSP Gx00 models, VSP Fx00 models, and VSP N series .

The report contains the following metrics:

- Model: The storage system type.
- Serial No: The serial number of the storage system.
- Total Efficiency Ratio: The ratio of the total saving effect achieved by accelerated compression, capacity saving (compression and deduplication), snapshot, and Dynamic Provisioning of storage systems, plus the start and end times for the calculation.



Note: Under certain circumstances, the Total Efficiency Ratio can display as 99999 (for example: before data has been written to newly-created pool or virtual volume).

- Data Reduction Ratio: The data reduction ratio before and after accelerated compression and capacity saving (compression and deduplication).
- Software Saving Ratio: The capacity reduction ratio before and after capacity saving.

- Software Compression Ratio: The capacity compression ratio before and after capacity saving.
- Software Deduplication Ratio: The capacity deduplication ratio before and after capacity saving.
- Software Pattern Matching Ratio: The capacity reduction ratio before and after pattern matching of capacity saving.
- FMD Saving Ratio: The capacity reduction ratio before and after accelerated compression.
- FMD Compression Ratio: The capacity compression ratio before and after accelerated compression.
- FMD Pattern Matching Ratio: The capacity reduction ratio before and after pattern matching of accelerated compression.
- Snapshot Efficiency Ratio: The efficiency ratio achieved by snapshot.
- Provisioning Efficiency (%): The efficiency ratio achieved by Dynamic Provisioning.
- Calculation Time: The start date/time and the end date/time for the calculation.

Capacity of Drive Types by Storage System report (HDP)

The **Capacity of Drive Types by Storage System** report shows the capacity usage of drive types for VSP 5000 series, VSP E series, VSP F series, VSP G series, VSP N series, VSP, and HUS VM series storage systems.

Use this report to monitor the enterprise-wide Hitachi Dynamic Provisioning (HDP) total pool capacity, used capacity, and free capacity details.

Model	Serial No	Flash HDP Total	Flash HDP Used	Flash HDP Free	Flash Parity Group	HDD HDP Total	HDD HDP Used	HDD HDP Free	HDD Parity Group
VSP F800	510025	-	-	-	-	8278.03 GB	2263.76 GB	6014.28 GB	-

The report contains the following metrics:

- HDP Total: The sum of the capacity of HDP pools in the storage system.
- HDP Used: The sum of the used capacity of HDP pools the storage system.
- HDP Free: The sum of the free capacity of HDP pools in the storage system.
- Parity Group Free: The sum of the free capacity of parity groups in the storage system .
- Flash HDP Total: The sum of the capacity of HDP pools with flash drive type in the storage system.
- Flash HDP Used: The sum of the used capacity of HDP pools with flash drive type in the storage system.
- Flash HDP Free: The sum of the free capacity of HDP pools with flash drive type in the storage system.
- Flash Parity Group Free: The sum of the free space of parity groups with flash drive type in the storage system.

- HDD HDP Total: The sum of the capacity of HDP pools with HDD drive type in the storage system.
- HDD HDP Used: The sum of the used capacity of HDP pools with HDD drive type in the storage system.
- HDD HDP Free: The sum of the free capacity of HDP pools with HDD drive type in the storage system.
- HDD Parity Group Free: The sum of the free space of parity groups with HDD drive type in the storage system.
- SCM HDP Total: The sum of the capacity of HDP pools with SCM drive type in the storage system.
- SCM HDP Used: The sum of the used capacity of HDP pools with SCM drive type in the storage system.
- SCM HDP Free: The sum of the free capacity of HDP pools with SCM drive type in the storage system.
- SCM HDP Parity Group Free: The sum of the free capacity of parity groups with SCM drive type in the storage system.
- FMC HDP Total: The sum of the capacity of HDP pools with FMC drive type in the storage system.
- FMC HDP Used: The sum of the used capacity of HDP pools with FMC drive type in the storage system.
- FMC HDP Free: The sum of the free capacity of HDP pools with FMC drive type in the storage system.
- FMC HDP Parity Group Free: The sum of the free capacity of parity groups with FMC drive type in the storage system.
- FMD HDP Total: The sum of the capacity of HDP pools with FMD drive type in the storage system.
- FMD HDP Used: The sum of the used capacity of HDP pools with FMD drive type in the storage system.
- FMD HDP Free: The sum of the free capacity of HDP pools with FMD drive type in the storage system.
- FMD Parity Group Free: The sum of the free space of parity groups with FMD drive type in the storage system.
- SSD HDP Total: The sum of the capacity of HDP pools with SSD drive type in the storage system.
- SSD HDP Used: The sum of the used capacity of HDP pools with SSD drive type in the storage system.
- SSD HDP Free: The sum of the free capacity of HDP pools with SSD drive type in the storage system.
- SSD Parity Group Free: The sum of the free space of parity groups with SSD drive type in the storage system.
- SAS HDP Total: The sum of the capacity of HDP pools with SAS drive type in the storage system.

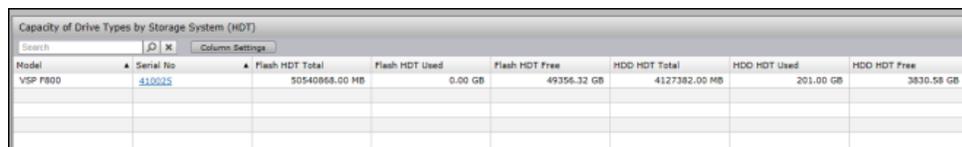
- SAS HDP Used: The sum of the used capacity of HDP pools with SAS drive type in the storage system.
- SAS HDP Free: The sum of the free capacity of HDP pools with SAS drive type in the storage system.
- SAS Parity Group Free: The sum of the free space of parity groups with SAS drive type in the storage system.
- SATA HDP Total: The sum of the capacity of HDP pools with SATA drive type in the storage system.
- SATA HDP Used: The sum of the used capacity of HDP pools with SATA drive type in the storage system.
- SATA HDP Free: The sum of the free capacity of HDP pools with SATA drive type in the storage system.
- SATA Parity Group Free: The sum of the free space of parity groups with SATA drive type in the storage system.

Click the link to the storage system serial number to open the Analyzer detail view UI in a separate browser. You can analyze the storage system performance reports on the Reports pane.

Capacity of Drive Types by Storage System report (HDT)

The **Capacity of Drive Types by Storage System** report shows the capacity usage of drive types for VSP 5000 series, VSP E series, VSP F series, VSP G series, VSP N series, VSP, and HUS VM series storage systems.

This report provides a breakdown of the enterprise-wide Hitachi Dynamic Tiering (HDT) total pool capacity, used capacity, and free capacity details. If an HDT pool has tiers from multiple drive types, then all the capacity in this pool is accounted for under the "Mixed" drive type.



Model	Serial No	Flash HDT Total	Flash HDT Used	Flash HDT Free	HDD HDT Total	HDD HDT Used	HDD HDT Free
VSP F800	413023	50540868.00 MB	0.00 GB	49356.32 GB	4127382.00 MB	201.00 GB	3830.58 GB

The report contains the following metrics for the supported drive types:

- HDT Total: The sum of the capacity of HDT tiers in a storage system.
- HDT Used: The sum of the used capacity of HDT tiers in a storage system.
- HDT Free: The sum of the free capacity of HDT tiers in a storage system.
- Flash HDT Total: The sum of the capacity of HDT tiers with flash drive type in a storage system.
- Flash HDT Used: The sum of the used capacity of HDT tiers with flash drive type in a storage system.
- Flash HDT Free: The sum of the free capacity of HDT tiers with flash drive type in a storage system.
- HDD HDT Total: The sum of the capacity of HDT tiers with HDD drive type in a storage system.

- HDD HDT Used: The sum of the used capacity of HDT tiers with HDD drive type in a storage system.
- HDD HDT Free: The sum of the free capacity of HDT tiers with HDD drive type in a storage system.
- FMC HDT Total: The sum of the capacity of HDT tiers with FMC drive type in a storage system.
- FMC HDT Used: The sum of the used capacity of HDT tiers with FMC drive type in a storage system.
- FMC HDT Free: The sum of the free capacity of HDT tiers with FMC drive type in a storage system.
- SCM HDT Total: The sum of the capacity of HDT tiers with SCM drive type in a storage system.
- SCM HDT Used: The sum of the used capacity of HDT tiers with SCM drive type in a storage system.
- SCM HDT Free: The sum of the free capacity of HDT tiers with SCM drive type in a storage system.
- FMD HDT Total: The sum of the capacity of HDT tiers with FMD drive type in a storage system.
- FMD HDT Used: The sum of the used capacity of HDT tiers with FMD drive type in a storage system.
- FMD HDT Free: The sum of the free capacity of HDT tiers with FMD drive type in a storage system.
- SSD HDT Total: The sum of the capacity of HDT tiers with SSD drive type in a storage system.
- SSD HDT Used: The sum of the used capacity of HDT tiers with SSD drive type in a storage system.
- SSD HDT Free: The sum of the free capacity of HDT tiers with SSD drive type in a storage system.
- SAS HDT Total: The sum of the capacity of HDT tiers with SAS drive type in a storage system.
- SAS HDT Used: The sum of the used capacity of HDT tiers with SAS drive type in a storage system.
- SAS HDT Free: The sum of the free capacity of HDT tiers with SAS drive type in a storage system.
- SATA HDT Total: The sum of the capacity of HDT tiers with SATA drive type in a storage system.
- SATA HDT Used: The sum of the used capacity of HDT tiers with SATA drive type in a storage system.
- SATA HDT Free: The sum of the free capacity of HDT tiers with SATA drive type in a storage system.

- External HDT Total: The sum of the capacity of HDT tiers with External drive type in a storage system. External storage consists of drives or other storage devices that are physically located in an externally connected storage system.
- External HDT Used: The sum of the used capacity of HDT tiers with External drive type in a storage system. External storage consists of drives or other storage devices that are physically located in an externally connected storage system.
- External HDT Free: The sum of the free capacity of HDT tiers with External drive type in a storage system.
- Mixed HDT Total: The sum of the capacity of HDT tiers with Mixed drive type in a storage system.
- Mixed HDT Used: The sum of the used capacity of HDT tiers with Mixed drive type in a storage system.
- Mixed HDT Free: The sum of the free capacity of HDT tiers with Mixed drive type in a storage system.

Click the link to the storage system serial number to open the Analyzer detail view UI in a separate browser. You can analyze the storage system performance reports on the Reports pane.

Capacity of Drive Types by Storage System report (HUS 100 Storage Systems)

The **Capacity of Drive Types by Storage System** report shows the capacity usage of drive types for HUS 100 series.

Use this report to monitor the Hitachi Dynamic Provisioning (HDP) total pool capacity, used capacity, and free capacity details for HUS100 storage systems.

Model	Serial No	Flash Thin Total	Flash Thin Used	Flash Thin Free	Flash Parity Group	HDD Thin Total	HDD Thin Used	HDD Thin Free	HDD Parity Group
HUS130	92110310	-	-	-	-	7280.00 GB	4651.00 GB	2629.00 GB	56.00 GB

The report contains the following metrics:

- Thin Total: The sum of the capacity of HDP pools with the target drive type in the storage system.
- Thin Used: The sum of the used capacity of HDP pools with the target drive type in the storage system.
- Thin Free: The sum of the free capacity of HDP pools with the target drive type in the storage system.
- Parity Group Free: The sum of the free space of parity groups with the target drive type in the storage system.
- Flash Thin Total: The sum of the capacity of HDP pools with Flash drive type in the storage system.
- Flash Thin Used: The sum of the used capacity of HDP pools with Flash drive type in the storage system.

- Flash Thin Free: The sum of the free capacity of HDP pools with Flash drive type in the storage system.
- Flash Parity Group Free: The sum of the free space of parity groups with Flash drive type in the storage system.
- HDD Thin Total: The sum of the capacity of HDP pools with HDD drive type in the storage system.
- HDD Thin Used: The sum of the used capacity of HDP pools with HDD drive type in the storage system.
- HDD Thin Free: The sum of the free capacity of HDP pools with HDD drive type in the storage system.
- HDD Parity Group Free: The sum of the free space of parity groups with HDD drive type in the storage system.
- FMD Thin Total: The sum of the capacity of HDP pools with FMD drive type in the storage system.
- FMD Thin Used: The sum of the used capacity of HDP pools with FMD drive type in the storage system.
- FMD Thin Free: The sum of the free capacity of HDP pools with FMD drive type in the storage system.
- FMD Parity Group Free: The sum of the free space of parity groups with FMD drive type in the storage system.
- SSD Thin Total: The sum of the capacity of HDP pools with SSD drive type in the storage system.
- SSD Thin Used: The sum of the used capacity of HDP pools with SSD drive type in the storage system.
- SSD Thin Free: The sum of the free capacity of HDP pools with SSD drive type in the storage system.
- SSD Parity Group Free: The sum of the free space of parity groups with SSD drive type in the storage system.
- SAS Thin Total: The sum of the capacity of HDP pools with SAS drive type in the storage system.
- SAS Thin Used: The sum of the used capacity of HDP pools with SAS drive type in the storage system.
- SAS Thin Free: The sum of the free capacity of HDP pools with SAS drive type in the storage system.
- SAS Parity Group Free: The sum of the free space of parity groups with SAS drive type in the storage system.

Click the link to the storage system serial number to open the Analyzer detail view UI in a separate browser. You can analyze the storage system performance reports on the Reports pane.

Capacity by Pool report

The **Capacity by Pool** report provides the breakdown of the total physical capacity, capacity savings, subscribed capacity and consumed capacity of a pool. Use this report to monitor the consumption and subscription rate of pools across all storage systems within your enterprise and determine if you are at risk of oversubscribing pools. You can track the capacity savings achieved by using data reduction technologies such as controller-based compression, deduplication, and accelerated compression.

Model	Serial No	Pool	Logical Capacity	% of Used Logical Capacity	Subscription %	Threshold Rate	Dedupe/Comp Capacity
VSP G1500	51457	51457.HSA_DDM_POOL_0as	2143.89 GB	95.00 %	0.00 %	100 %	0.00 GB
VSP G1500	51457	51457.emoc7(11)	495.88 GB	18.00 %	8347.53 %	80 %	202.90 GB
VSP G1500	51457	51457.hfvm-seale-normalbo	45.90 GB	0.00 %	12.96 %	80 %	0.00 GB
VSP G1500	51457	51457.seethereid(10)	9311.61 GB	0.00 %	0.00 %	80 %	0.00 GB
VSP G1500	51457	51457.hac-jethare-endname	2196.67 GB	1.00 %	22.78 %	93 %	0.00 GB

This report applies to VSP 5000 series, VSP E series, VSP Gx00 models, VSP Fx00 models, and VSP N series:

The report contains the following metrics:

- **Model:** The storage system type to which the pool belongs.
- **Serial No:** The serial number of the storage system to which the pool belongs.
- **Pool:** The pool ID.
- **Pool Type:** The type of the pool.
- **Logical Capacity:** The logical capacity of the pool.
- **Logical Used Capacity:** The logical capacity used in a pool.
- **% of Used Logical Capacity:** The percentage of used logical capacity in a pool.
- **Physical Capacity:** The physical capacity of the pool.
- **Physical Used Capacity:** The physical capacity used in a pool.
- **Subscribed Capacity:** The sum of the capacity of DP volumes in a pool.
- **Subscription %:** The subscription percentage of a pool.
- **Threshold Rate:** The threshold set for capacity use in a storage system.
- **Demand:** The sum of the free capacity of allocated DP volumes in a pool.
- **Auto Expansion:** The capability of expanding the pool automatically.
- **Data Volumes Used Capacity:** Total used capacity of data volumes (DP-VOLs).
- **Data Reduction Capacity:** The storage capacity savings with the controller-based capacity saving and accelerated compression functions in a pool.
- **Data Reduction Ratio:** The storage capacity savings ratio with the controller-based capacity saving and accelerated compression functions in a pool.
- **Deduplication Capacity:** The saved capacity with the capacity saving (deduplication) function in a pool.
- **Deduplication Ratio:** The saving ratio with the capacity saving (deduplication) function in a pool.
- **Compression Capacity:** The saved capacity with the capacity saving (compression) function in a pool.

- Compression Ratio: The saving ratio with the capacity saving (compression) function in a pool.
- Dedupe/Comp Capacity: The saved capacity with the capacity saving function.
- Dedupe/Comp Ratio: The saving capacity ratio with the capacity saving function.
- Pattern Matching: The reclaimed capacity with the capacity saving function.
- System Data: The system data capacity (metadata capacity plus garbage data capacity) saved with the capacity saving function.
- Preprocessed Data: The preprocessed data capacity with the capacity saving function.
- FMC Logical total: The total capacity with the accelerated compression function.
- FMC Logical Used: The total used capacity with the accelerated compression function.
- FMC Physical Total: The total physical capacity with the accelerated compression function.
- FMC Physical Used: The used physical capacity with the accelerated compression function.
- FMC Compression Capacity: The compressed capacity with the accelerated compression function.
- FMC Compression Ratio: The saved capacity ratio with the accelerated compression function.
- FMC Pattern Matching: The reclaimed capacity with the accelerated compression function.

Click the links to the storage system serial number or pool ID to open the Analyzer detail view UI in a separate browser. You can analyze the storage system performance reports and the pool performance reports on the Reports pane.

Capacity by Pool report (HUS 100 Storage Systems)

The **Capacity by Pool** report for HUS 100 Storage Systems provides the breakdown of the total physical capacity, subscribed capacity, and consumed capacity of a pool. Use this report to monitor the consumption, and subscription rate of pools that belong to the HUS 100 series and determine if you are at risk of oversubscribing pools.



Model	Serial No	Pool	Capacity	Capacity Used %
HUS130	92110310	92110310-0	5461.00 GB	80 %
HUS130	92110310	92110310-5	1819.00 GB	15 %

The report displays the following metrics:

- Model: The storage system type to which the pool belongs.
- Serial No: The serial number of the storage system to which the pool belongs.
- Pool: The pool ID.
- Drive Type: The type and speed of the physical drives.
- Capacity: The capacity of the pool.

- Used Capacity: The capacity used in a pool.
- Capacity Used %: The percentage of capacity used in a pool.
- Subscribed Capacity: The sum of the capacity of DP volumes in a pool.
- Demand: The sum of the free capacity of allocated DP volumes in a pool.

Click the links to storage system serial number or pool ID to open the Analyzer detail view UI in a separate browser. You can analyze the storage system performance reports and the pool performance reports on the Reports pane.

Total Efficiency by Pool report

The **Total Efficiency by Pool** report provides a summary of the capacity saved through data reduction features such as controller-based compression, deduplication, and accelerated compression at the pool level.

The report contains the following metrics:

- Model: The storage system type to which the pool belongs.
- Serial No: The serial number of the storage system.
- Pool: The pool ID.
- Pool Type: The type of pool.
- Total Efficiency Ratio: The ratio of the total saving effect achieved by accelerated compression, capacity saving (compression and deduplication), snapshot, and Dynamic Provisioning that belong to storage systems, followed by the start date and time and the end date and time for the calculation.



Note: Under certain circumstances, the Total Efficiency Ratio can display as 99999 (for example: before data has been written to newly-created pool or virtual volume).

- Data Reduction Ratio: The data reduction ratio before and after accelerated compression and capacity saving (compression and deduplication).
- Software Saving Ratio: The capacity reduction ratio before and after capacity saving.
- Software Compression Ratio: The capacity compression ratio before and after capacity saving.
- Software Deduplication Ratio: The capacity deduplication ratio before and after capacity saving.
- Software Pattern Matching Ratio: The capacity reduction ratio before and after pattern matching of capacity saving.
- FMD Saving Ratio: The capacity reduction ratio before and after accelerated compression.
- FMD Compression Ratio: The capacity compression ratio before and after accelerated compression.
- FMD Pattern Matching Ratio: The capacity reduction before and after pattern matching of accelerated compression.
- Snapshot Efficiency Ratio: The efficiency ratio achieved by snapshot.

- Provisioning Efficiency (%): The efficiency ratio achieved by Dynamic Provisioning.
- Calculation Time: The start date/time and the end date/time for the calculation.

Capacity by Consumer report

The **Capacity by Consumer** report provides the current snapshot of the storage capacity used by the consumers. The report provides a breakdown of used capacity, free capacity, and other usage metrics for DP volumes and non-DP volumes. Use this report to monitor the capacity consumption of resources allocated to a consumer group.

Capacity by Consumer				
Search		Column Settings		
Consumer ▲	Grade ▲	Thin Used	Thin Free	Thick
Banking and Fina...	Platinum	520.40 GB	179.60 GB	48.40 GB
IT Help Desk Ser...	Gold	324.00 GB	64.00 GB	16.00 GB
HR Logistics and ...	Gold	128.50 GB	96.50 GB	20.50 GB
ERP Services	Bronze	64.00 GB	128.00 GB	32.00 GB
#Unassigned Re...	-	3052.60 GB	757217.50 GB	848001.70 GB

The report contains the following metrics:

- Consumer: The name of the consumer.
- Grade: The grade of the consumer.
- Thin Used: The sum of the used capacity of DP volumes assigned to the consumer.
- Thin Free: The sum of the free capacity of DP volumes assigned to the consumer.
- Thick: The sum of the capacity of non-DP volumes assigned to the consumer.

Capacity by Consumer in the Past 6 Months report

The **Capacity by Consumer in the Past 6 Months** report displays the storage capacity used by the consumers in the past 6 months. The report provides a breakdown of used capacity, free capacity, and other usage metrics for DP volumes and non-DP volumes. Use this report to monitor the capacity consumption of resources allocated to a consumer group.

Capacity by Consumer in the Past 6 Months				
Search		Column Settings		
Consumer ▲	Grade ▲	Thin Used	Thin Free	Thick
Banking and Fina...	Platinum	1280.40 GB	3268.60 GB	248.40 GB
IT Help Desk Ser...	Gold	512.00 GB	1048.00 GB	960.00 GB
HR Logistics and ...	Gold	268.50 GB	892.50 GB	480.50 GB
ERP Services	Bronze	128.00 GB	512.00 GB	532.00 GB
#Unassigned Re...	-	128643.20 GB	963848.50 GB	649632.00 GB

The report contains the following metrics:

- Consumer: The name of the consumer.
- Grade: The grade of the consumer.

- Thin Used: The sum of the used capacity of DP volumes assigned to the consumer.
- Thin Free: The sum of the free capacity of DP volumes assigned to the consumer.
- Thick: The sum of the capacity of non-DP volumes assigned to the consumer.

Capacity by Consumer report (HUS 100 Storage Systems)

The **Capacity by Consumer (HUS 100 Storage Systems)** report provides details about the storage capacity used by the consumers. The report provides a breakdown of used capacity, free capacity, and other usage metrics for DP volumes, and non-DP volumes. Use this report to monitor the capacity consumption of resources allocated to a consumer group.

Capacity by Consumer (HUS 100 Storage Systems)				
Search		Column Settings		
Consumer	Grade	Thin Used	Thin Free	Thick
Banking and Fina...	Platinum	520.40 GB	179.60 GB	48.40 GB
IT Help Desk Ser...	Gold	324.00 GB	64.00 GB	16.00 GB
HR Logistics and ...	Gold	128.50 GB	96.50 GB	20.50 GB
ERP Services	Bronze	64.00 GB	128.00 GB	32.00 GB
#Unassigned Re...	-	4594.92 GB	1548.48 GB	5409.82 GB

The report contains the following metrics:

- Consumer: The name of the consumer.
- Grade: The grade of the consumer.
- Thin Used: The sum of the used capacity of DP volumes assigned to the consumer.
- Thin Free: The sum of the free capacity of DP volumes assigned to the consumer.
- Thick: The sum of the capacity of non-DP volumes assigned to the consumer.

Capacity by Consumer in the Past 6 Months report (HUS 100 Storage Systems)

The **Capacity by Consumer in the Past 6 Months (HUS 100 Storage Systems)** report provides details about the storage capacity used by consumer in the past 6 months. The report provides a breakdown of used capacity, free capacity, and other usage metrics for DP volumes and non-DP volumes. Use this report to monitor the capacity consumption of resources allocated to a consumer group.

Capacity by Consumer in the Past 6 Months (HUS 100 Storage Systems)				
Search <input type="text"/>		Column Settings		
Consumer ▲	Grade ▲	Thin Used	Thin Free	Thick
Banking and Fina...	Platinum	1280.40 GB	3268.60 GB	248.40 GB
IT Help Desk Ser...	Gold	512.00 GB	1048.00 GB	960.00 GB
HR Logistics and ...	Gold	268.50 GB	892.50 GB	480.50 GB
ERP Services	Bronze	128.00 GB	512.00 GB	532.00 GB
#Unassigned Re...	-	128643.20 GB	963848.50 GB	649632.00 GB

The report contains the following metrics:

- Consumer: The name of the consumer.
- Grade: The grade of the consumer.
- Thin Used: The sum of the used capacity of DP volumes assigned to the consumer.
- Thin Free: The sum of the free capacity of DP volumes assigned to the consumer.
- Thick: The sum of the capacity of non-DP volumes assigned to the consumer.

Chapter 12: Performance analytics and best practices

Understanding utilization

The single, most important cause of performance problems is an end-to-end (E2E) data path to a resource where the throughput capacity has been exceeded. In this case, infrastructure that supports the large-scale movement of data is bound to experience issues with throughput. To fulfill the promise of no single point of failure, all resources in the data path must have sufficient reserve capacity to carry increased load under failover conditions.

Performance analysis frequently begins with a determination of whether any E2E data path resources are overloaded. By using E2E data analysis in Hitachi Ops Center Analyzer (Analyzer) you can determine the particular areas of concern.

To understand how utilization is measured, consider the categories capacity and workload. Capacity refers to the percentage of physical space occupied across storage resources. Using percentages allows you measure the space occupied on a disk or a parity group. Not so straightforward is the category of workload, which can be thought of as the percentage of time it takes for the system to do an operation (for example, read or write). Another way of expressing workload is to measure the percentage of how busy a resource, such as a processor or cache, is.

These resources include those on the host and storage side, as well as the fabric. Because the whole system and its parts are affected by workload, the time and space aspect of measuring utilization can be combined to focus on capacity throughput, which is used to measure port and path utilization. Throughput capacity combines both considerations by calculating the percentage of I/O operations.

When monitoring infrastructure resource workload with Analyzer, use the following metrics to measure utilization:

- **IOPS (I/O per second):** the number of operations
- **MBps (MB per second) :** the amount of data transferred
- **Response time (read/write operations):** the sum of service and wait time.

Optimizing online versus batch workloads

Optimizing utilization of a storage resource for an online workload and a batch workload is mutually exclusive. Consequently, online and batch workloads generally do not share the same storage access resource at the same time.

The performance for online workloads is monitored on a minute by minute basis. Performance alerts us to such problems as, at the most pressing, lagging response times, and while less immediate but just as significant, inadequate reserves to support processing during failures.

For batch workloads, the average performance over time is monitored. Since high utilization is a design goal, the metrics for these values are not a cause for concern: maximizing utilization per resource maximizes throughput per resource. Consequently, high response time is not central to monitoring batch workloads. Instead, monitoring batch workloads becomes a question of whether there is adequate capacity to complete processing within the batch workload window even after a component failure occurs.

Utilization metrics

To evaluate performance, track the percentage of utilization of active resources and throughput.

You should baseline your I/O profile when all systems are operating in a normal state. The values shown in the Normal Value column are planning estimates.

Metric Name	Name in detail view	Normal Value (in milliseconds)	Performance risk value (in milliseconds)
Analyzer			
Read response time	Average disk reads/ second	1 to 10	Greater than 10
Write response time	Average disk reads/ second	1 to 3	Greater than 3

Storage-area network considerations

Storage-area network (SAN) components are also part of the end-to-end data path analysis. Areas of concerns that need to be addressed could include inter-switch links (ISLs) and latency issues due to throughput. The switch configuration might not be in line with the desired infrastructure design, so we need to maintain the correct firmware and make sure we do not have any issues with SFP.

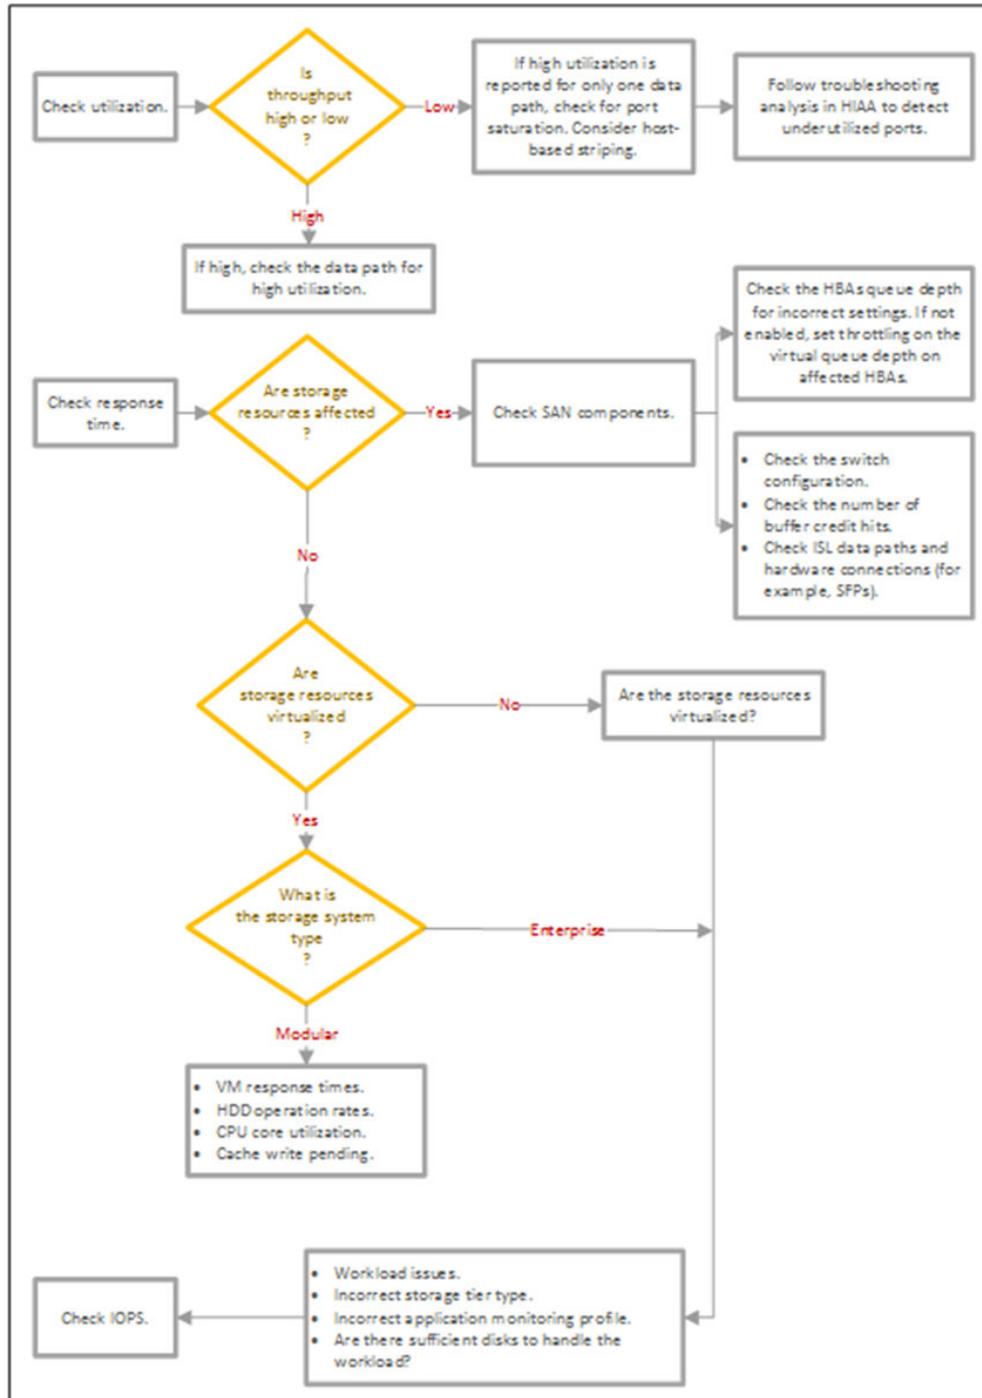
For example, a 48-port board has 2 ASICs sharing ports:

- If communication is between ports on the same ASIC port, then you see a 700-nanosecond latency.
- If communication is between ports of different ASICs, or between different boards, traffic goes through the back panel, through core communication boards, and then at the end of the transfer goes through three ASICs which could take 2 microseconds.

Regarding throughput, each port can transfer the maximum capability of the ASIC (i.e., 16 Gbps with Gen5, 32 Gbps with newer switch architectures. Most switches can sustain the number of IOPS necessary to offer maximum throughput with a normal frame size of 2048 bytes. This might not be true with smaller frames.

Performance analysis

Use the following decision-making workflow to check performance in your infrastructure.



Performance analysis example

Use the following workflow as an example of the troubleshooting sequence for analyzing performance by metrics. In this scenario, performance degradation has caused a critical alert to be sent about a platinum grade consumer in Ops Center Analyzer.

Analyze response time

1. Check performance of the affected storage resources:
 - a. See performance metrics for parity groups, MPB, I/O paths, response time, and cache write pending.
 - b. If the storage system is modular, see performance metrics for UVM/VM response time, HDD operation rates, CPU core busy, cache write pending.
2. Consider the following components:
 - a. HBA queue depth
 - b. HBA settings
 - c. Virtual queue depth settings on the ESX host
3. Check the following switch components:
 - a. Switch configuration
 - b. Buffer-to-buffer credit hits
 - c. ISL links
 - d. Hardware failures (for example, an SFP module might not be properly connected)

Analyze IOPS

1. Check for workload issues.
 - Is the storage tier type incorrect?
 - Is the profile for this application correct?
 - Is there enough disk space to handle the workload?
2. Check the performance for Parity Groups, MPBs, I/O paths.

Check utilization

If throughput is high:

1. Check all associated data paths for high utilization.
2. Check the CHA>ESX host paths.

If throughput is low:

- Check the host paths for high utilization.
- Check for port saturation on the host.
- Consider host-based striping.

Performance troubleshooting best practices

Use the following best practices for infrastructure reporting.

Report timing

Most performance issues are analyzed using one-minute data intervals. However, performance problems do occur requiring shorter interval analysis. Analysis of one-minute data is generally limited to one- or two-day durations. Short intervals avoid muting peaks by averaging performance metric values.

Troubleshoot high response times

When you monitor your infrastructure resources, the most significant metric to watch in the online transactions is the I/O rate. The application processes a large number of transactions when the I/O rates are higher. Maintain healthy response times in an OLTP environment, which mostly generates random access I/O. The read I/O response times should be higher than the write response times. Use the following guidelines when considering response time thresholds:

- This threshold depends on the application requirements and the SLA.
- Since the LUN response time has a direct impact on applications, this indicator should be monitored on key LUNs to determine deltas as loads increase.
- Look for the worst performing LUNs and correlate with the host disk.

Maximum recommended array group utilization

For online workloads:

- 50% during normal operations
- Utilization reserves are required to accommodate failure

For batch workloads:

- As high as possible, because batch metric is typically given in elapsed time
- Expect maximums of 70-80% (depending on the burst profile of initiator)
- Average utilization over time remaining in the batch window should not exceed 50%

VSP processor maximum planned utilization with capacity reserves to accommodate failure

- VSP Virtual Storage Director (VSD) manages a specific list of LDEVs.
- VSPs are redundant MPB pairs with 1:1 failover.
- VSD is monitored by analyzing MPB utilization, as shown in the following table.

	Maximum Recommended Planned MPB Utilization	Notes
Online	40%	1 minute interval average (includes reserve for failover)
Online	70%	1 minute interval average
Batch	40%	1 minute interval average (includes reserve for failover)
Batch	80%	1 minute interval average

VSD cache considerations

- The cache for each CLPR is allocated to the respective VSDs.
- The cache allocation for each VSD is initially uniform among VSDs, but can be dynamically reallocated in response to changing load conditions.
- Inflow control due to high write pending levels is local to the VSD/CLPR pair having a high write pending level.
- Accelerated de-staging due to high write pending levels is a system-wide activity.
- Up to 30% write pending is considered normal.
- Frequent or sustained increases to 40% deserve attention.
- Frequent or sustained increases to 50% deserve prompt attention.
- At 70%, emergency de-stage/inflow control is invoked.

Appendix A: Definition file templates

Ops Center Analyzer provides users with definition file templates that can be customized to run actions based on specific events, along with email notifications.

Setting event actions

Setting event actions allows you to run a batch file for running an event action when an Ops Center Analyzer event is registered. You can enable automatic notifications when an error is detected by Ops Center Analyzer by defining commands to run when an event is registered in the batch file for running an event action.

Defining a file for running an event action

To run an event action, create a file that defines the commands to run when a specific event is registered.

You can specify any file name. The files must have a file extension of `.sh`.

In the file for running an event action, you can view information about the event that triggered an event action through environment variables. The following table shows the environment variables that you can specify in a file for running an event action.

Variable name	Description
ANALYTICS_SOURCE	Device name
ANALYTICS_DEVICE	Device type
ANALYTICS_DESCRIPTION	Message
ANALYTICS_CATEGORY	Category
ANALYTICS_SEVERITY	Level
ANALYTICS_DATE	Registration date
ANALYTICS_EVENTID	Event ID
ANALYTICS_GROUPS <i>nnn</i>	Group name. Specify 000 for <i>nnn</i> .
ANALYTICS_NODEID	Node ID
ANALYTICS_COMPONENTID	Component ID

Variable name	Description
ANALYTICS_PERFCOMPONENTID	Performance
ANALYTICS_NAME	Name of a host where Ops Center Analyzer is running

Example of specifying an environment variable in Linux: `${ANALYTICS_SOURCE}`

Event action definition file format

The following describes the format used to enter data and the settings to specify in `EventAction.properties`. The timing when the specified definitions are applied is also described.

Format

specified-key-name=specified-value

File

`EventAction.properties`

Save the file in UTF-8 format. When you save the file, prevent a BOM (byte order mark) from being added to the file.

Folder

Ops-Center-Analyzer-installation-folder/Analytics/conf

The default installation folder is `/opt/hitachi`.

Update frequency

Ops Center Analyzer rereads the definition file if more than 5 minutes have elapsed since the previous reading before executing the event action.

Content to specify

Specify each key name and value on a single line. When you create the event action definition file, the following rules apply:

- A line starting with `#` is treated as a comment line
- Blank lines are ignored
- The entered values are case-sensitive
- If the same key is specified more than once in the same file, the last specification is valid
- To specify a backslash (`\`), enter `\\`
- To display `%`, specify `%%`

- To specify an apostrophe ('), enter \'
- To display a double quotation mark ("), specify \"
- Specify only an absolute path. A path specified in an environment variable cannot be set for the path

Setting description

Key name	Setting description	Specifiable values	Default value	Optional or required
EventAction.cmd	Specifies the absolute path of the batch file for running an event action.	ASCII characters and characters in 260 bytes, excluding control characters. Note: Spaces are excluded.	Null character	Optional If this key is omitted, no event action will be run. This key also enables or disables the function.
EventAction.maxCount	Specifies the maximum number of event actions that can be run simultaneously.	1 to 100	10	Optional If this key is omitted, the default value will be used.
EventAction.timeOut	Specifies the timeout time for event actions (in milliseconds).	1 to 3,600,000	300,000	Optional If this key is omitted, the default value will be used.

Format of the email template definition file

To create emails more efficiently, create an email template definition file and then use an email template to create emails from the **Execute Action** window. You can also create emails that contain Ops Center Analyzer-managed resource data. For example, you can create an email template that provides information on how to troubleshoot a resource failure. When a failure occurs, email templates allow you to create emails from the **Execute Action** window to quickly notify the system administrator.

Prerequisite

- To send email messages from Ops Center Analyzer, configure the email server for Ops Center Analyzer.

Format

specified-key-name=specified-value

File

- Use any file.
- Save the file in UTF-8 format.
- A maximum of 1,000 files can be set in Ops Center Analyzer. Files are loaded in alphabetical order by file name, and any files after the 1,000th file are not loaded.

Folder

*Ops-Center-Analyzer-installation-folder/Analytics/conf/
template/mail*

The default installation folder is `/opt/hitachi`.

Update frequency

When Ops Center Analyzer is started or the `reloadtemplate` command is run.

Content to specify

Specify each key name and value on a single line. When you create the email template definition file, the following rules apply:

- A line starting with `#` is treated as a comment line
- Blank lines are ignored
- The entered values are case-sensitive
- If the same key is specified more than once in the same file, the last specification is valid
- To display `\`, specify `\\`
- To display `%`, specify `%%`
- If the filter condition `SE.template.filter.xxxxxxx.string` is specified more than once, settings will be displayed when all of the conditions are met
- If you specify "LFCR" for the setting value, it displays in a new line in a preview window

Setting description

Key name	Setting description	Specifiable values	Default value	Optional or required
SE.template.name.string	Specify the action name.	Values of no more than 127 bytes that do not include control characters	Null character	Required If this key is omitted, the processing to read files will fail. The name of this key is the same as the name specified in a command template definition file.
SE.template.description.string	Specify a description of the action.	Values of no more than 255 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used.
SE.mail.template.title.string	Specify the subject of the email template.	Values of no more than 255 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used.
SE.mail.template.body.string	Specify the body of the email template.	Values of no more than 4,096 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used.
SE.mail.template.address.string	Specify the address of the email template	Values of no more than 255 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used.

Key name	Setting description	Specifiable values	Default value	Optional or required
SE.template.filter.resourceName.string	Specify conditions for the names of the resources that are starting points and that display in the action list during resource selection. ¹	Values of no more than 255 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used, and the key will not be used as a filter condition.
SE.template.filter.resourceType.string	Specify conditions for the types of resource that are starting points and that display in the action list during resource selection. ¹	Values of no more than 32 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used, and the key will not be used as a filter condition.
SE.template.filter.vmHostname.string	Specify conditions for the virtual machine names that display in the action list during resource selection. ¹	Values of no more than 64 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used, and the key will not be used as a filter condition.
SE.template.filter.ipaddress.string	Specify conditions for the IP addresses that display in the action list during resource selection. ¹	Values of no more than 255 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used, and the key will not be used as a filter condition.

Key name	Setting description	Specifiable values	Default value	Optional or required
SE.template.filter.upperResourceName.string	Specify conditions for the names of higher-level resources of a starting point that display in the action list during resource selection. ¹	Values of no more than 512 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used, and the key will not be used as a filter condition.
SE.template.filter.upperResourceType.string	Specify conditions for the types of higher-level resources of a starting point that display in the action list during resource selection. ¹	Values of no more than 32 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value will be used, and the key will not be used as a filter condition.
¹ Settings display only when the Execute Action window is called from a resource that matches the specified conditions.				

By using variables, you can set information about a selected resource as the value of a setting.

Variable name	Description	Remarks
%ANALYTICS_RESOURCE_NAME%	Name of the selected resource	Not applicable
%ANALYTICS_UPPERRESOURCE_NAME%	Name of the higher-level resource of the selected resource	Not applicable
%ANALYTICS_IPADDRESS%	IP address	Not applicable
%ANALYTICS_VIRTUALMACHINE_NAME%	Name of the virtual host	Displays only when the resource is a virtual machine.

Variable name	Description	Remarks
%ANALYTICS_RESOURCE TYPE%	Resource type	Not applicable
%ANALYTICS_UPPERRES OURCETYPE%	Type of the higher-level resource	Not applicable

If no value is set for the selected resource, a null character displays.

To display information about virtual hosts and IP addresses, VMware Tools must be installed on virtual hosts.



Note:

The email template definition file restricts the maximum number of characters that can be displayed in the email editor.

If the maximum number of specifiable characters is exceeded in an email template's address, subject, and body, clicking Launch Mailer in the **Execute Action** dialog box might not start the email editor normally.

If the email editor does not start normally, manually start the email editor, and then copy the contents of the preview to use them.

The maximum number of specifiable characters depends on the web browser. As a guideline, the maximum number of characters for web browsers is 2059 for the Firefox browser, and 200 for the Internet Explorer 11 browser.

Command template definition files formats

If you create a command template definition file, use the **Execute Action** window to run commands of other products, user programs, and resources on the Ops Center Analyzer management server. If you want to run more than one command at the same time, you can create a batch or shell command.

Format

specified-key-name=specified-value

File

- Specify any file name and file extension.
- Save the file in UTF-8 format.
- The maximum number of files that can be set in Ops Center Analyzer (including the number of email template definition files) is 1,000. Files load in alphabetical order by file name, and any files after the 1,000th file are not loaded.

Folder

*Ops-Center-Analyzer-installation-folder/Analytics/conf/
template/command*

The default installation folder is `/opt/hitachi`.

Update frequency

When Ops Center Analyzer is started or the `reloadtemplate` command is run.

Content to specify

Specify each key name and value on a single line. The following rules apply when specifying settings in a command template definition file:

- A line starting with `#` is treated as a comment line.
- Blank lines are ignored.
- The entered values are case-sensitive.
- If you specify an invalid value, the default value is used.
- If you specify the same key more than once in the same file, the last specification is valid.
- To display `\`, specify `\\`.
- To display `%`, specify `%%`.
- If you specify the filter condition `SE.template.filter.xxxxxxx.string` more than once, settings display when all of the conditions are met.

Setting descriptions

Key name	Setting description	Specifiable values	Default value	Optional or required
<code>SE.template.name.string</code>	Specify the action name.	Values of no more than 127 bytes that do not include control characters	This setting has no default value, because specifying this setting is required.	Required
<code>SE.template.description.string</code>	Specify a description of the action.	Values of no more than 255 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value is used.

Key name	Setting description	Specifiable values	Default value	Optional or required
<code>SE.cmd.template.cmdName.string</code>	Specify the name of the command to run by specifying the absolute path to the command. A command specified by its relative path might fail.	Values of no more than 255 bytes that do not include control characters Do not specify the following strings: & ; \$ > < ` ! ../ .. \ () { }	This setting has no default value, because specifying this setting is required.	Required
<code>SE.cmd.template.cmdArgs.string</code>	Specify arguments for the command to run.	Values of no more than 4,096 bytes that do not include control characters Do not specify the following strings: & ; \$ > < ` ! ../ .. \ () { }	Null character	Optional If this key is omitted, the default value is used.
<code>SE.cmd.template.timeout.num</code>	Specify the timeout period for the command to run (in milliseconds).	1 to 2,147,483,647	30,000	Optional If this key is omitted, the default value is used.
<code>SE.template.filter.resourceName.string</code>	Specify conditions for the names of the resources that are starting points and that display in the action list during resource selection. ¹	Values of no more than 255 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value is used.

Key name	Setting description	Specifiable values	Default value	Optional or required
SE.template.filter.resourceType.string	Specify conditions for the types of resources that are starting points and that display in the action list during resource selection. ¹	Values of no more than 32 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value is used.
SE.template.filter.vmHostname.string	Specify conditions for the virtual machine names that display in the action list during resource selection.*	Values of no more than 64 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value is used.
SE.template.filter.ipaddress.string	Specify conditions for the IP addresses that display in the action list during resource selection. ¹	Values of no more than 255 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value is used.
SE.template.filter.upperResourceName.string	Specify conditions for the names of higher-level resources of a starting point that display in the action list during resource selection. ¹	Values of no more than 512 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value is used.
SE.template.filter.upperResourceType.string	Specify conditions for the types of higher-level resources of a starting point that display in the action list during resource selection. ¹	Values of no more than 32 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value is used.

Key name	Setting description	Specifiable values	Default value	Optional or required
SE.template.filter.MultipleResources.boolean	Specify whether to execute for multiple resources.	true or false	false	Optional If this key is omitted, the default value is used.
SE.cmd.template.usage.string	Specify how to use the command.	Values of no more than 4,096 bytes that do not include control characters	Null character	Optional If this key is omitted, the default value is used.
* Settings display only when the Execute Action window is called from a resource that matches the specified conditions.				

By using variables, you can set information about a selected resource as the value of a setting.

The following table lists the variables you can use.

Variable name	Description	Remarks
%ANALYTICS_RESOURCE_NAME%	Name of the selected resource	N/A
%ANALYTICS_UPPERRESOURCE_NAME%	Name of the higher-level resource of the selected resource	N/A
%ANALYTICS_IPADDRESS%	IP address	N/A
%ANALYTICS_VIRTUALMACHINE_NAME%	Name of the virtual host	Displays only when the resource is a virtual machine.
%ANALYTICS_RESOURCE_TYPE%	Resource type	N/A
%ANALYTICS_UPPERRESOURCE_TYPE%	Type of the higher-level resource	N/A

If no value is set for the selected resource, a null character displays.

To display information about virtual hosts and IP addresses, VMware Tools must be installed on virtual hosts.

Definition example where the selected resource names are specified for command arguments

```
SE.template.name.string=001_task-execution
SE.template.description.string=Executes the scheduled tasks
SE.cmd.template.cmdName.string=/root/commands/taskA.sh
SE.cmd.template.cmdArgs.string=%ANALYTICS_RESOURCENAME%
```

Appendix B: Troubleshooting

The common performance problems and the possible solutions are described in this appendix.

Solving performance problems

The common performance problems and the possible solutions are described as follows. The possible causes and solutions are intended to provide guidance, and might not satisfy your business process performance requirements.

The following table lists the commonly observed storage-related problems and possible solutions.

Bottleneck metrics	Root cause and possible solutions
Parity Group utilization	<ul style="list-style-type: none">▪ Root cause: The usage rate of the Parity Group increases because of the following possible causes:<ul style="list-style-type: none">▪ Some volumes might be under heavy load.▪ Volumes (logical resources) might belong to the same Parity Group (physical resource) which might cause resource contention issues in the shared infrastructure.▪ Possible solutions:<ul style="list-style-type: none">• Consider moving some volumes to another Parity Group with a lower usage rate or higher performance.• Consider increasing the number of drives (by concatenating Parity Groups).• To manage a Parity Group that is part of a pool, consider adding another Parity Group to the pool.

Bottleneck metrics	Root cause and possible solutions
MPB utilization	<ul style="list-style-type: none"> ▪ Root cause: The usage rate of the MP Blade (average usage rate of the MP cores in the MP Blade) increases because of an increased load. Too many busy resources such as internal volumes, external volumes, or journal groups accessing the same MP Blade might cause performance degradation. ▪ Possible solutions: <ul style="list-style-type: none"> • Consider allocating the busy resources (internal volumes, external volumes, or journal groups) to another MP Blade (changing the MP Blade ownership). • Limit I/O to the volumes using I/O controls. • Increase the size of cache memory allocated to the MP Blade.
Port utilization	<ul style="list-style-type: none"> ▪ Root cause: The usage rate of the port (amount of data forwarded by the port divided by the amount of data that can be forwarded by the port) increases because of a number of volumes accessing the same port. ▪ Possible solutions: Consider allocating some volumes (or host groups) to a different port. Note: When the connected port is changed, the host might need to be restarted.

Bottleneck metrics	Root cause and possible solutions
Cache utilization	<ul style="list-style-type: none"> ▪ Root cause: <p>Out of the total cache memory allocated to the CLPR, the percentage occupied by the data waiting to be written to the drive increases. The cache write pending rate increases because of the following possible causes:</p> <ul style="list-style-type: none"> • The usage rate of the drive might be high, delaying write processing to the drive. • The usage rate of the processors might be high, delaying write processing to the drive. • The capacity of the installed cache memory might be insufficient. ▪ Possible solutions: <ul style="list-style-type: none"> • Consider allocating some volumes to another cache partition. • Consider increasing the cache memory. • Limit I/O to the volumes using I/O controls. • Consider allocating the busy resources (internal volumes, external volumes, or journal groups) to another MP Blade (changing the MP Blade ownership).

The following table lists the commonly observed server-related problems and possible solutions.

Bottleneck metrics	Root cause and possible solutions
CPU utilization	<ul style="list-style-type: none"> ▪ Root cause: CPU bottlenecks occur when several VMs run on the same physical machine, and end up sharing the same CPU. If the VMs (logical resources) share the same CPU (physical resource) and if one of the VMs utilizes the CPU more than the others in the shared infrastructure, the total efficiency of the CPU is degraded and the CPU utilization rate increases. The CPU could become saturated with requests because of resource contention issues. ▪ Possible solutions: Consider moving the VMs to another server.
Memory utilization	<ul style="list-style-type: none"> ▪ Root cause: Memory bottlenecks occur when several VMs (logical resources) share the available memory (physical resources) which might result in the performance degradation of the physical memory. ▪ Possible solutions: Consider allocating additional physical memory, or moving the VMs to another server.

Troubleshooting granular data collection error codes

If an error occurs during granular data collection, determine the cause of the error using one of the following methods.

- Check the message display on the Events tab in Ops Center Analyzer.

If data collection does not start because the script corresponding to an action does not exist, the message returns a code that indicates the cause of the error.

Example:

Message: Command finished. (*action name*, command return code: *return code*)

Refer to the following table to identify the cause of the error and the solution.

- Check the result in the log file.

If the return code 0 is displayed on the Events tab, check the log file to see the results of the granular data collection.

Message: KNAQ03626-I command finished. (exit code: *exit code*)



Note: The message ID related to granular data collection is KNAQ03626-I.

If an action ends successfully, the message displays exit code 0. If an exit code other than 0 is displayed, refer to the following table to identify the cause of the error and the solution.

Return code	Root cause and solutions
10-21	<p>An invalid value is specified for an option.</p> <p>Refer to the log and specify an appropriate value for the option that caused the error.</p>
22-23	<p>The value specified for the <code>-startTime</code> or <code>-timezone</code> option is invalid.</p> <p>Specify appropriate values for these options.</p>
24	<p>The date and time specified for the <code>-startTime</code> option is earlier than the current date and time.</p> <p>Specify an appropriate value for this option.</p>
28	<p>The target Agent host could not be identified.</p> <ul style="list-style-type: none"> ▪ Make sure that the file <code>storage_agent_map.txt</code> exists in the same directory as this command. ▪ Confirm that the lines corresponding to the model and serial numbers specified in the <code>-model</code> and <code>-serialNumber</code> options are coded in the file <code>storage_agent_map.txt</code>.
29	<p>The port number defined in the file <code>storage_agent_map.txt</code> is invalid.</p> <p>Make sure that a numeric value is defined for the port number.</p>
30	<p>A connection to the Agent host could not be established.</p> <ul style="list-style-type: none"> ▪ Make sure that the IP address of the Agent host specified in the file <code>storage_agent_map.txt</code> is correct. ▪ Make sure that Agent host is running. ▪ Confirm that communication can be established between the Ops Center Analyzer server and the Agent host. ▪ Make sure that the initial settings for this command are completed.
40-44	<p>An internal error occurred.</p> <ul style="list-style-type: none"> ▪ Make sure that IP address of the Agent host specified in the file <code>storage_agent_map.txt</code> is correct. ▪ Make sure that the Agent host is running. ▪ Confirm that communication can be established between the Ops Center Analyzer server and the Agent host. ▪ Make sure that the initial settings of this command are completed.
45	<p>The storage information could not be obtained.</p> <ul style="list-style-type: none"> ▪ Make sure that the Tuning Manager Web service is running on the Agent host.

Return code	Root cause and solutions
	<ul style="list-style-type: none"> ▪ Make sure that the port number of the Agent host specified in the file <code>storage_agent_map.txt</code> matches the port number used by the Tuning Manager Web service.
46	<p>The Agent instance that is to monitor the target storage system does not exist.</p> <ul style="list-style-type: none"> ▪ Make sure that the storage system specified by the <code>-model</code> and <code>-serialNumber</code> options is set to be monitored on the corresponding Agent host. ▪ Check that the PD records have been already collected on the instance that is to monitor the specified storage system.
50	<p>The directory that stores the collection results exists as a file.</p> <ul style="list-style-type: none"> ▪ Refer to the log file and delete or rename the file whose path is shown for the <code>RESULT_DIR</code> parameter. ▪ Alternatively, change the path name specified for the <code>-pathLabel</code> option.
51	<p>The directory that stores the execution results exists as a file.</p> <ul style="list-style-type: none"> ▪ Refer to the log file and delete or rename the file <code>model-name_serial-number</code> whose path is shown for the <code>RESULT_DIR</code> parameter. ▪ Alternatively, change the path name specified for the <code>-pathLabel</code> option.
52	<p>The directory that stores the execution results could not be created.</p> <p>Refer to the log file and check the permissions for the directory shown for the <code>RESULT_DIR</code> parameter.</p>
53-54	<p>An internal error occurred. A temporary file could not be created or deleted.</p> <p>Refer to the log file and check the permissions for the directory shown for the <code>RESULT_DIR</code> parameter.</p>
60-62	<p>An internal error occurred. An operation could not be performed for a file or directory on the Agent host.</p> <p>Delete the file in the following directory on the Agent host:</p> <pre data-bbox="488 1503 1354 1562">/opt/jplpc/agtd/agent/instance-name/raidperf_ldevlist/hiaa.conf</pre> <pre data-bbox="488 1587 1227 1612">/opt/jplpc/agtd/agent/instance-name/secdata_rm</pre> <p>For <i>instance-name</i>, specify the name of the instance that monitors the specified storage system.</p>
70-72	<p>An internal error occurred. The RAID Agent command could not be executed normally on the Agent host.</p> <ul style="list-style-type: none"> ▪ Make sure that the Agent host is running and that communication is available between the Ops Center Analyzer server and the Agent host.

Return code	Root cause and solutions
	<ul style="list-style-type: none"> ▪ Make sure that the Agent version is 8.5.1 or later. ▪ Part of the pre-transfer execution result might remain in the following directory on the Agent host. Because this directory is initialized the next time the command is run, obtain the file beforehand if necessary: <code>/opt/jplpc/agtd/agent/instance-name/secdata_rm</code> For <i>instance-name</i>, specify the name of the instance that monitors the specified storage system.
73	<p>The execution result file could not be transferred from the Agent host. Part of the pre-transfer execution result might remain in the following directory on the Agent host.</p> <p>Because this directory is initialized the next time the command is run, obtain the file beforehand, if necessary.</p>
80	<p>An internal error occurred.</p> <p>Make sure that Agent host is running and that communication is available between the Ops Center Analyzer server and the Agent host.</p> <p><code>/opt/jplpc/agtd/agent/instance-name/secdata_rm</code></p> <p>For <i>instance-name</i>, specify the name of the instance that monitors the specified storage system.</p>
81	<p>The data collection processing is being executed for the same storage by the same command.</p> <p>After the data collection processing is finished, run the command again.</p> <p>If this error occurs but no data collection processing is running, delete the following file on the Agent host:</p> <p><code>/opt/jplpc/agtd/agent/instance-name/raidperf_ldevlist_hiaa.conf</code></p> <p>For <i>instance-name</i>, specify the name of the instance that monitors the specified storage system.</p>
90	<p>The installation path of Ops Center Analyzer was not found.</p> <p>Check the status of the installation.</p>
91	<p>The schedule could not be set according to the time specified in the <code>-startTime</code> option.</p> <ul style="list-style-type: none"> ▪ Check the specification for this option. ▪ Make sure that the <code>atd</code> service is running on the Ops Center Analyzer server.

Return code	Root cause and solutions
92	<p>The command required for specifying the start of the collection time (the <code>-startTime</code> option) was not found.</p> <p>Install the <code>at</code> command on the Ops Center Analyzer server.</p>

Index

Special Characters

assignment rule

- changing the order of priority (resource) 111
- changing the order of priority (user resource) 105
- creating (resource) 110
- creating (user resource) 104
- executing (resource) 111
- executing (user resource) 105

A

- action template 54, 56
- analytics 20, 69, 75
- analyze 69, 75
- Analyzer detail view 10.8.1 20, 69, 75
- analyzing
 - bottleneck 46, 52
- authentication
 - external, groups 119
 - external, users 118

B

- batch file for event action execution 163
- bottleneck
 - analyzing 46
 - determining cause 51

C

- Capacity of VMware Datastores report 141
- Capacity of VMware Datastores with Usage Exceeding 80% report 141
- capacity reports
 - drive types 145
 - HUS 100 storage systems 143
- condition profile
 - creating 115
- configuration report 31
- configuring
 - mail server 114
- consumer

consumer (*continued*)

- creating 110, 112
- grade 124
- status 124
- summary 124
- Consumer - Critical report 125
- consumer capacity 154
- consumers
 - settings 109
- creating
 - consumer 110, 112
- critical reports
 - Consumer - Critical 125
- custom report 31
- custom reports
 - line graphs 31
 - table report 31

D

- dashboard
 - overview 22
- datastores
 - Capacity of VMware Datastores report 141
 - Capacity of VMware Datastores with Usage Exceeding 80% report 141
- default reports
 - Event Trends 26
 - Resource Events 27
 - System Resource Status 25, 123
 - System Status Summary 24
- definition 163
- definition file template 170
- device mapping 65

E

- email notification
 - condition profile (creating) 115
 - settings 114
- email notifications
 - email address (settings to enable or disable) 115

- email template definition file 165
- event action 163
- event action definition file 164
- Event Trends report 26
- events
 - overview 27
- Execute Action
 - Granular Data Collection 57

F

- format 164, 165, 170

G

- granular data collection 56
- Granular Data Collection
 - collecting granular data 57
 - error codes and solutions 180
 - guidelines 57
 - options 57
 - output 57
 - prerequisites 57
 - troubleshooting 180

H

- Hitachi Ops Center Analyzer 10.8.1 20, 69, 75
- Hitachi Ops Center Automator 20
- Hitachi Performance Analytics 20, 69, 75
- hypervisor 135

I

- I/O control
 - clear settings 89
 - optimizing I/O performance 86
 - set IO control 88
 - SLO 86

M

- mail server
 - configuring 114
- monitoring environment setup 20
- monitoring settings
 - dynamic threshold (overview) 94
 - static threshold (overview) 99

O

- Ops Center Analyzer
 - features 14

- Ops Center Analyzer (*continued*)
 - logging on 18
 - overview 13
- Ops Center Analyzer 10.8.1 20, 69, 75
- Ops Center Analyzer detail view 10.8.1 20, 69, 75
- Ops Center AutomatorAutomation Director 20
- Ops Center Automatorservices 56

P

- Performance Analytics 20, 69, 75
- predict 69, 75
- predictive analytics
 - generating a risk report 73
 - licenses 71
 - report definitions 70
 - risk profile 70
- predictive risk profile 70
- predictive risk report definitions 70

R

- report definition 69, 75
- reports
 - Capacity of VMware Datastores 141
 - Capacity of VMware Datastores with Usage Exceeding 80% 141
 - Consumer - Critical 125
 - Event Trends 26
 - Resource Events 27
 - storage systems 137
 - switches 136
 - System Resource Status 25, 123
 - system resources 138
 - System Status Summary 24
 - Total Hosts 134
 - Total VMs 129
 - VM CPU Ready 128
 - VM Disk Latency 129
 - VM NIC Dropped 128
- resource
 - assignment rule (changing the order of priority) 111
 - assignment rule (creating) 110
 - assignment rule (executing) 111
- Resource Events report 27
- resources
 - monitoring settings (overview) 91
- risk profile 69, 75
- risk profiles 72
- risk report definitions 72
- risk reporting 69, 75

S

- second-level data 56
- security 120
- service template 56
- settings
 - consumers 109
 - email notification 114
 - monitoring resources (overview) 91
- storage capacity 142
- system resource
 - setting thresholds 105
- System Resource Status report 25, 123
- System Status Summary
 - infrastructure resources 122
 - Storage Resources 121
 - System Resources 122
 - user resources
 - host 25, 122
 - VM 25, 122
 - volume 25, 122
 - Volumes 121
- System Status Summary report 24

VMs (continued)

- Total VMs report 129

T

- threshold
 - settings (system resource) 105
 - settings (user resource) 103
- topology view 65
- Total Hosts report 134
- Total VMs report 129
- trend chart 69, 75

U

- user resource
 - assignment rule (changing the order of priority) 105
 - assignment rule (creating) 104
 - assignment rule (executing) 105
 - setting thresholds 103

V

- VM CPU Ready report 128
- VM Disk Latency report 129
- VM NIC Dropped report 128
- VM reports
 - VM CPU Ready 128
 - VM Disk Latency 129
 - VM NIC Dropped 128
- VMs

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact