



Hewlett Packard
Enterprise

HPE Cloud Optimizer

Software Version: 3.00
Linux operating system

PDF version of the Online Help

Document Release Date: January 2016
Software Release Date: January 2016

Legal Notices

Warranty

The only warranties for Hewlett-Packard Development Company, L.P. products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2012 - 2016 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® is a U.S. registered trademark of the Microsoft group of companies.

Java is a registered trademark of Oracle and/or its affiliates.

Acknowledgements

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)

This product includes software written by Tim Hudson (tjh@cryptsoft.com)

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<https://softwaresupport.hp.com>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<https://hpp12.passport.hp.com/hppcf/createuser.do>

Or click the **the Register** link at the top of the HP Software Support page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at: <https://softwaresupport.hp.com>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<https://hpp12.passport.hp.com/hppcf/createuser.do>

To find more information about access levels, go to:

<https://softwaresupport.hp.com/web/softwaresupport/access-levels>

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is

<http://h20230.www2.hp.com/sc/solutions/index.jsp>

Contents

Chapter 1: Introduction to HPE Cloud Optimizer	10
Chapter 2: About the HPE Cloud Optimizer Console	11
Chapter 3: Configuring HPE Cloud Optimizer Using Settings	13
All	14
Adding a Datasource	16
Removing a Datasource	17
Adding vSphere Datasource	17
Adding Microsoft Hyper-V Datasource	20
Proxy and Remote WMI Configuration	22
Adding KVM Datasources	22
Adding XEN Datasources	24
Adding OpenStack Datasources	25
Adding OneView Datasource	30
Adding Amazon Web Services Datasource	30
Adding Physical Server as a Datasource	32
Monitoring Physical Servers using HPE Cloud Optimizer Features	34
Monitoring the Health of a Datasource	36
Restarting Collection	37
License	37
Manage License	38
Monitoring License	38
Integrating HPE Cloud Optimizer with Other HPE Products	39
Integrating HPE Cloud Optimizer with HPE Cloud Service Automation (HPE CSA)	40
Integrating HPE Cloud Optimizer with CSA from HPE Cloud Optimizer Console	41
Monitoring the Infrastructure as a Cloud Service Automation (CSA) user	42
Configuring CSA for HPE Cloud Optimizer Single Sign-On Utility	43
Integrating BSM with HPE Cloud Optimizer	44
Integrating HPE Cloud Optimizer with HPE Operations Manager	45
Integrating HPE Cloud Optimizer with HPE Operations Manager i (OMi)	48
Viewing Performance Graphs in HPE OMi	53
Customizing the Severity of Alerts	55
Integrating HPE Cloud Optimizer with HPE OneView	55
Using HPE Cloud Optimizer features with HPE OneView Integration	57
Unified Enclosure View	58
Using Unified Enclosure View	60
Working with Business Groups	67
Creating a Business Grouping	68
Using the Business Grouping Creation Wizard	69

Using Business Groups with other HPE Cloud Optimizer Features	72
Resource Meter	76
Business Metric Analyzer	79
Configuring HPE Cloud Optimizer	83
Configuring Session Timeout Interval	85
Configuring Secure Client Connect Timeout Interval	86
Configuring Java Virtual Machine Memory	86
Configuring HPE Cloud Optimizer to use Public Key Infrastructure Authentication	87
Disabling Certificate-based Authentication	89
Configuring LDAP	89
Accessing HPE Cloud Optimizer User Interface	94
Daily Maintenance	94
Time Zones	95
 Chapter 4: Dashboard	 96
Using the HPE Cloud Optimizer Dashboard	96
Monitoring Environment Performance	97
Monitoring Environment Capacity	99
Viewing Forecast Data	100
Viewing Alerts	101
Customizing the Dashboard View	101
 Chapter 5: Treemap	 103
Treemap View	103
Using Treemap View	104
Resources Pane	106
Using the Options Pane	108
Drill Down to Resources	111
Cross launch to actual domain from Treemap	112
Searching Resources	112
Filtering Data	113
Micro Charts	113
Viewing Micro Charts	114
Use Case: Viewing Data in the Treemap	114
 Chapter 6: Performance	 116
Performance Graphs	116
Graphs Overview	117
Icon Glossary	117
Workbench Page	118
Configuration Items (CI)	118

Favorites	119
Performance Pane	119
Options for Drawn Graphs	120
Table Graph Window	123
Using Table Highlights	124
Using Table Filters	125
Graph Export Dialog Box	126
Date Range Panel	126
Drawing Graphs	128
Saving as Favorite	129
Deleting a Favorite	129
Drawn Graphs Functions	130
Reports Overview	131
Viewing Reports	132
Types of Reports	133
Using HPE Cloud Optimizer Reports for vSphere	133
Datacenter	134
Cluster	135
ESX/ESXi Host	136
Resource Pool	137
Datastore	138
VM	138
Using HPE Cloud Optimizer Reports for Hyper-V	139
Host Cluster	139
Host	140
Datastore	140
VM	141
Using HPE Cloud Optimizer Reports for OpenStack	141
Cloud	141
Tenant	142
Host	142
VM	143
Using HPE Cloud Optimizer Reports for KVM/Xen	143
Host	143
VM	144
Using HPE Cloud Optimizer reports for HPE OneView	144
Enclosures	144
ServerHardwares or Blade Servers	145
VMware Cluster	146
Filter	146
 Chapter 7: Alerts	 148
Using Alert Messages	150
Alert View	151
Conditions Used to Generate the Alert Messages in VMware	152

Conditions Used to Generate the Alert Messages in Hyper-V	156
Conditions Used to Generate the Alert Messages in KVM	158
Monitoring Alerts on the HPOM Console	159
Sensitivity	160
Configuring Sensitivity	161
Suppressing Alerts	162
Monitoring the State of Virtual Machines	166
Managing vCenter Events and Alarms	168
 Chapter 8: Capacity	 170
Personas	170
Using Capacity	170
Overview	171
Summary for Datacenter	176
Viewing Snapshots List	176
Summary for Cluster	178
Summary for Host	178
Summary for VM	181
Summary for Datastore	184
CPU Sizing	186
Memory Sizing	188
Configuring Memory Sizing	189
View Idle VMs	190
View Powered off VMs	191
CPU and Memory Reservation	192
Use Case: Using Capacity Usage Trends and Optimization Option to Manage the Resource Utilization in your Environment	193
Using Placement	196
Forecast	198
Forecasting CPU Utilization	201
Forecasting Memory Utilization	202
Forecasting Disk Utilization	202
Using Forecast to plan for resources	203
Modeler	204
Scenarios	208
Find Capacity	209
Configuring Capacity Head Room Settings	212
Configure the Head Room Values in pvcd.PO namespace	212
Using the Command Line	213
 Chapter 9: Troubleshooting HPE Cloud Optimizer	 214
Frequently Asked Questions	229

Chapter 10: Real Time Guest OS Drilldown for Advance Troubleshooting	238
Guest OS DrillDown User Interface	239
Use Case: Using Guest OS Drilldown to Troubleshoot the Resource Utilization in your Environment	246
Chapter 11: HPE Cloud Optimizer API Reference	248
Using REST API for Placement	248
Using REST APIs for Optimization Recommendations	270
Datacenter	276
To get information of all the Datacenters	276
To get information of a specific Datacenter	277
To get sizing information of all the VMs that are part of a specific Datacenter	278
To get sizing information of a specific VM in a Datacenter filtered based on name and sizing parameters	280
To get sizing information of all the datastores that are part of a specific Datacenter	283
Cluster	284
To get information of all clusters	284
To get information of a specific cluster	286
To get sizing information of all the VMs in a specific cluster	287
To get sizing information of a specific VM in a cluster filtered based on name and sizing parameters	289
To get sizing information of all the datastores that are part of a specific cluster	291
Host	292
To get information of all the Hosts	292
To get information of a specific Host	294
To get sizing information of all the VMs in a specific Host	297
To get sizing information of a specific VM in a Host filtered based on name and sizing parameters	299
To get datastore information and sizing details of all the datastores in a specific Host	302
Virtual Machines	303
To get information about all the VMs along with sizing details	303
To get information of a specific VM	305
Datastores	308
To get datastore information along with sizing details	308
To get information of a specific datastore	310
To get information about all the VMs on a specific datastore	311
To get information of a specific VM in a datastore filtered on the basis of Status	313
Using REST API to view Powered Off VMs	314
Using REST APIs to view Idle VMs	318
Using REST APIs to view CPU, Memory, and Storage allocation of VMs	322
Using REST APIs to view Snapshots	333
Appendix	336

Send Documentation Feedback	337
-----------------------------------	-----

Chapter 1: Introduction to HPE Cloud Optimizer

HPE Cloud Optimizer is a web-based analysis and visualization tool that analyzes performance trends of elements in virtualized environments. It enables virtualization monitoring by providing an overview of the environment, near-real-time and historical data analysis and triaging using an interactive dashboard. It also enables monitoring for cloud, hypervisor, non-hypervisor or physical environments. HPE Cloud Optimizer helps you visualize performance data for elements in the context of each other to rapidly analyze bottlenecks. HPE Cloud Optimizer provides performance monitoring, graphing, and reporting in a single interface.

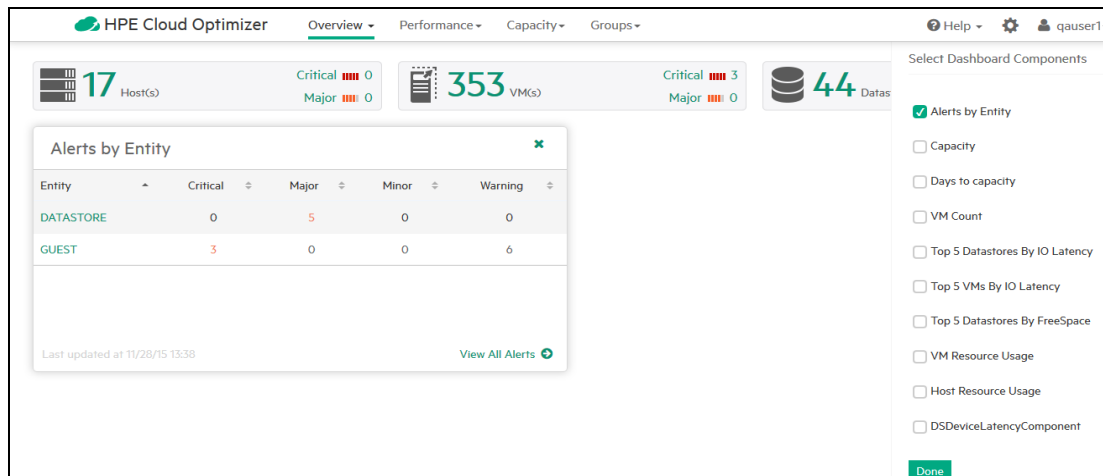
Key features of HPE Cloud Optimizer are as follows:

- Monitor the health and performance of the virtualized environment from Dashboard view.
- Triage analysis with the Workbench and capability to trend server utilization across days, weeks, and a month.
- Proactive monitoring of the entities that may cause disruptions in the environment. Use Treemap to identify and view the problems at a glance.
- Analyze the capacity, usage, and allocation trends for various resources in a virtualized environment.
- Right sizing recommendation based on historical resource utilization and reclaiming unused resources.
- Analyze the impact of increasing the workload.
- Determine the impact of adding or deleting resources in your environment.
- Estimate the future capacity utilization in order to proactively plan your hardware requirements.
- Ability to group VMs, hosts, and datastores into Business Groups for collective analysis in Treemap, Workbench, and Forecast.

Note: HPE Cloud Optimizer supports the VMware vCenter Server versions 5.0, 5.1, 5.5, and 6.0

Chapter 2: About the HPE Cloud Optimizer Optimizer Console

The HPE Cloud Optimizer console consists of the following features as shown in the diagram:



Feature	Description
Overview	<p>Quick overview of performance, capacity, usage, and forecast details of entities in your environment. To view the data, navigate to the following:</p> <ul style="list-style-type: none"> Dashboard - Overview of the performance and usage of entities in your environment Treemap - Visual representation of the performance of resources monitored Unified Enclosure View - Representation of the virtual infrastructure overlaid on the actual physical layout.
Performance	<p>View the performance data of the entities. To view the data, navigate to the following:</p> <ul style="list-style-type: none"> Workbench - Draw graphs and generate reports for performance data. Alerts - Monitor alerts in the environment. Filter - Filters shown data in the Treemap.

Feature	Description
Capacity	<p>View the capacity and usage data of entities. To view the data, navigate to the following:</p> <ul style="list-style-type: none"> • Overview - Capacity usage trends of entities in environment. • Forecast - View forecast on resource utilization. • Modeler - View impact of adding resources in the environment. • Placement - Recommends the distribution of new resources in the environment.
Group	<p>This tab is available only when you are working with business groups. To view the data, navigate to the following:</p> <ul style="list-style-type: none"> • Business Metric Analyzer - Allows you to use Business Groups to effectively plan and manage capacity allocation. • Resource Meter - Allows you to calculate the billing cost.
Settings	<p>To administer the following HPE Cloud Optimizer functions:</p> <ul style="list-style-type: none"> • Datasources - Add datasources to start monitoring the resources in the environment. • Integrations - Integrate HPE Cloud Optimizer with other HPE products. • License Management - Information about HPE Cloud Optimizer licenses. <p>For more information, see "Configuring HPE Cloud Optimizer Using Settings" on page 13</p>
Help	<p>Access help on the product.</p> <ul style="list-style-type: none"> • Help for HPE Cloud Optimizer - Opens the Online Help documentation. • Documentation Library - Consists of links to the complete HPE Cloud Optimizer Documentation Set. It also contains many additional resources to help you plan, deploy, and best use HPE Cloud Optimizer. • Troubleshooting Tips - View troubleshooting tips for different domains. • About HPE Cloud Optimizer - Know product specific details.

Chapter 3: Configuring HPE Cloud Optimizer Using Settings

As a HPE Cloud Optimizer Administrator, you can use Settings to configure HPE Cloud Optimizer.

What tasks can you perform using Settings?

The following table lists the tasks you can perform as an administrator of HPE Cloud Optimizer:

Configuration tasks	Action
Adding or Removing Datasources	Use the Collection & Analysis option to add or remove datasources.
Configuring HPE Cloud Optimizer Integration	Use the Integration option to integrate and use HPE Cloud Optimizer with other HPE products - HPE Business Service Management, HPE Performance Manager, and HPE Cloud Service Automation. For more information, see Integrating with Other HPE Products .
Managing HPE Cloud Optimizer Licenses	Use the License option to manage your HPE Cloud Optimizer licenses. You can also activate your evaluation license and import permanent licenses using this option. For more information, see License .
Working with Business groups	Use the Business Group option to create, deploy, and delete business groups.

Note: The Virtual Appliance Management Interface (VAMI) is provided by VMware and hence is not available in the localized languages.

Note: Make sure Tomcat is configured with HTTPS for secure communication. For more information on configuring secure communication, see *Configuring Secure Communication with Tomcat Server* section in *HPE Cloud Optimizer Installation Guide*.

Using the Settings page

Click the options to know more.

All ₁	Collection & Analysis ₂	License ₃	Integration ₄	Business Group ₅
------------------	------------------------------------	----------------------	--------------------------	-----------------------------

	Option	Description
1	All	The All option displays all the available options in the Settings page.
2	Collection & Analysis	The Collection & Analysis option helps you select a domain, add a datasource to be monitored, remove any added datasource.
3	License	The License option provides information about your HPE Cloud Optimizer licenses.
4	Integration	The Integration option helps you integrate HPE Cloud Optimizer with other HPE products.
5	Business Group	The Business Group option in HPE Cloud Optimizer allows you to customize the operational view of your virtual infrastructure, based on your business applications and services.

Note: You can use the Search functionality available in the Setting page to filter data based on your requirement.

Navigation

Click the **Settings** icon  on the upper-right corner.

All

The **All** option displays all the available options in the **Settings** page. This option is selected by default when you click Settings.

From the All option, you can access the following features:


Collection & Analysis

License

Integration

Business Groups

Note: The **Launch Cloud Optimizer Virtual Appliance Management** link is used to change settings like Time zone, Network Settings for the HPE Cloud Optimizer Virtual Appliance.

The  is used for self monitoring of HPE Cloud Optimizer and will take you to [Real Time Guest OS Drilldown for Advanced Troubleshooting](#).

Collection and Analysis

Using the **Collection & Analysis** option you can select a domain, add a datasource to be monitored, remove any added datasource, and also see the list of datasources that are monitored by HPE Cloud Optimizer. After adding a datasource, the status and details of the datasource appear in the **Status of Datasources** list at the bottom of the page.

Following are the tasks that you can perform from the Collection & Analysis option:

- [Adding a Datasource](#)
- [Removing a Datasource](#)
- [Monitoring the health of a Datasource](#)

The following table lists the elements available in the Status of Datasources List:

Name	Description
IP/URL/Host Name	Lists the IP addresses, URLs, or host names of all the datasources added to HPE Cloud Optimizer for monitoring.
Domain	Domain of the datasource added for monitoring.
User Name	User name corresponding to the added datasource.
Licensed Instance Count	Number of instances associated with each monitored datasource. Total instance count is the total number of VMs and Hosts being monitored.
Status	Shows the current status of each connection. For the list of available

Name	Description
	status messages, see "Data Collection Status Messages" .
Last Collection Time	Shows the date and client time zone against each datasource, when data was last collected. Note: This field does not show any value till the first data collection is complete.
Actions	Includes Restart or Delete buttons. You can restart or delete a datasource.

Note:

- HPE Cloud Optimizer datasource collection works only when sum of instances from all datasources is less than or equal to the maximum instance capacity of the license. If the sum of the instances exceeds the maximum instance capacity of the license, then connection for all the datasources fails. For example, if you have installed a permanent license for 100 instances and try to connect two hosts with 60 instances each, then the collection for both the hosts fails.
- You can configure the Collection Interval for all the domains except for Hyper-V. Set the collection interval to 15 minutes for domains where collection takes more than 5 minutes to complete. By default, the collection interval is 5 minutes.

Adding a Datasource

To start monitoring the utilization of resources in your virtualized environment, you must add datasources to HPE Cloud Optimizer. When you add the datasource to be monitored, HPE Cloud Optimizer starts monitoring the resources and displays the relevant data in the Treemap, Graphs, and Reports. Using this data, you can troubleshoot performance issues in your environment. Based on this data, you can plan your resource utilization properly.

- [VMware vCenter](#)
- [Hyper-V](#)
- [OneView](#)
- [Physical Server](#)
- [Xen](#)
- [OpenStack](#)

- [KVM](#)
- [AWS](#)

Removing a Datasource

When you want to stop monitoring a datasource in your virtualized environment, you can remove the datasource from HPE Cloud Optimizer. After removing the datasource, HPE Cloud Optimizer stops monitoring that datasource and does not display any data corresponding to it.

Note: To remove remote collectors such as Hyper-V, run the following command:

```
pvconfig -dt -ty HYPERV -t <target name>
```

To remove a datasource from the list of Datasources Status, follow these steps:

1. Select the datasource you want to remove from the Status of Datasource list.
You can select multiple connections by pressing the `Ctrl` key and clicking the required connections.
2. Click **Delete** from Actions.
The system displays a dialog box asking for confirmation to remove the datasource.
3. Click **OK** to remove the datasource.
The system displays a confirmation message and no longer monitors the removed datasource. The datasource is also removed from the Status of Datasources list.

Note: When you remove a datasource from HPE Cloud Optimizer, Treemap and Workbench might display data corresponding to the deleted VMs for the next 2 data collection cycles (10 minutes).

Adding vSphere Datasource

HPE Cloud Optimizer supports the VMware vCenter Server versions 5.0, 5.1, 5.5, and 6.0.

Prerequisites

- The user must have the following roles in addition to the Read-only role:
 - **Browse Datastore** role available under **Datastore**
 - **Validate session** role available under **Sessions**
- You must **Enable Statistics** on the VMware vCenter server. VMware vCenter server

must have statistics level 2 enabled for 5 minutes and 30 minutes data logging.

To set the Statistics Level in the VMware vCenter server, follow these steps:

- a. In the vSphere Client, click **Administration** -> **vCenter Server Settings**.
- b. In the vCenter Server Settings window, click **Statistics**.

The Statistics Interval page displays the time interval after which the vCenter Server statistics will be saved, the time duration for which the statistics will be saved and the statistics level.

- c. Click **Edit**.
- d. In the Edit Statistics Interval window, set the following values:

Statistics Level:

For the statistics level that you select, the Edit Statistics Interval window displays the type of statistics which will be collected for that level.

Note: You must set the minimum statistics level as 2.

For more information, see the *VMware documentation*.

- The 443 (HTTPS) port should be open to allow HPE Cloud Optimizer collect data.
- In vSphere 5.5 Update 2d or higher and vSphere 6.0, a limit is set to the number of entities that are included in a database query. The limit protects the vCenter Server database from receiving large queries. The default value is 64. You must change the set limitation.

To change the limitation for performance charts you must perform two tasks:

- Edit the `config.vpxd.stats.maxQueryMetrics` key in the advanced settings of vCenter Server, or vCenter Server Appliance.
- Edit the `web.xml` file.

To change the `config.vpxd.stats.maxQueryMetrics` Key:

- a. Log on to vSphere Web or Desktop Client.
- b. Select **Administration > vCenter Server Settings** to display the vCenter Server Settings dialog box.
- c. On the **Settings** tab, click **Advanced Settings**.
- d. On the Advanced vCenter Server Setting page, click **Edit**.
- e. Navigate to the `config.vpxd.stats.maxQueryMetrics` key and change the value.

Note: If this is the first time you edit the key, you cannot see it in the list of advanced vCenter Server settings.

- f. Enter the name of the key and its value, and then click **Add**.

The default value is 64. To disable the limit, set a value to -1.

To edit the `web.xml` file for **vCenter Server** and **vCenter Server Appliance**

- a. Navigate to the `web.xml` file.
 - For vCenter Server Appliance, open `/usr/lib/vmware-perfcharts/tc-instance/webapps/statsreport/WEB-INF`
 - For vCenter Server, open
`C:\ProgramData\VMware\vCenterServer\data\perfcharts\tc-instance\webapps\statsreport\WEB-INF`
- b. Open the `web.xml` file in a text editor.
- c. Find the `maxQuerySize` parameter and change the value in the `<param-value>` tag. To disable the limit, set a value 0:

```
<context-param>
<description>Specify the maximum query size (number of metrics)for a
single report.
Non-positive values are ignored.</description>
<param-name>maxQuerySize</param-name>
<param-value>64</param-value>
</context-param>
```

- d. Save and close the `web.xml` file.
- e. In the vSphere Web Client, restart the VMware Performance Chart Service.

Caution: If you increase the value of the `config.vpxd.stats.maxQueryMetrics` setting, or if you disable its limit, you might send larger queries to the vCenter Server database. This might cause the following problems:

- Deterioration in the general responsiveness of the vCenter Server
- Increased response time to display some of the charts

For more information, see the *VMware documentation*.

To add a vSphere datasource, follow these steps:

1. Click **Collection & Analysis** option in the Settings page.
2. Click the **Add Datasource** link.
3. Select the **vSphere** option from the Domain drop-down.

4. Type the IP address or host name of the datasource in the **IP/URL/Host Name** field.

Note: If you add a VMware vCenter Server twice (with FQDN¹ and without FQDN), HPE Cloud Optimizer does not validate the VMware vCenter Server and adds it again.

5. Type the user name corresponding to the specified datasource in the **User Name** field.
6. Type the password corresponding to the specified datasource in the **Password** field.
7. Click **Test Connection**, if you want to validate the credentials provided.
The system validates the credentials and displays a message stating whether the credentials are valid or not. The system also validates if the connection is successful or not.
8. Click **Add**.

The Status of Datasources list displays the added datasources.

Adding Microsoft Hyper-V Datasource

Prerequisites

- Microsoft .NET Framework 4
- Windows 2012 R2 or Windows 2008 R2 SP1 system as proxy collector to collect data from Hyper-V
- Port used for accessing HPE Cloud Optimizer (8081/8444) should be open in the firewall
- Port between HPE Cloud Optimizer and Windows proxy system (383) should be open on both sides
- Port between Windows proxy system and Hyper-V host (135) should be open on Hyper-V host
- Microsoft Hyper-V 2008 R2 SP1 or Hyper-V Windows 2012 R2

To add a Hyper-V datasource, follow these steps:

1. Click **Collection & Analysis** option in the Settings page.
2. click the **Add Datasource** link.
3. Select **Hyper-V** from the Domain drop-down list.
4. Click **Download Collector** link to download the vPVWinVirtCollector.zip file.

¹Fully Qualified Domain Name

Extract the `vPVWinVirtCollector.zip` file on the proxy host.

The **vPVWinVirtCollector.zip** package contains the collector (**vPVWinVirtCollector.exe**), configuration file (**vPVWinVirtCollector.properties**), batch file (**Install.bat**), and dlls (**system.management.automation.dll**, **JsonSharp.dll**) with configurable parameters to install the collection service.

- The downloaded package is pre-configured with IP of the HPE Cloud Optimizer Instance, to which the collected data is reported. For example:
`vPVURL=http://<vPV host name or IP address>:383/pvcd.pvc.rrListener.JSONRemoteRequestHandler`
- The configuration file (**vPVWinVirtCollector.properties**) and log file is available at the same location where you have extracted the `vPVWinVirtCollector.zip` package.

5. Run the `Install.bat` as a domain user to start the collection. A pop-up window appears.

You must have read access to Windows Management Instrumentation (WMI) and appropriate access to the Performance Data from Targets. For more information on Proxy and Remote WMI Configuration, see ["Proxy and Remote WMI Configuration" on the next page](#)

6. Enter the domain user name and password in the window. The Installer installs the HPE Cloud Optimizer Collector Service and starts the collection service for targets mentioned in the properties file.
7. Open the `vPVWinVirtCollector.properties` file. Update the value for Targets. The value for target can be a list of **HYPERV** servers to be monitored separated by semicolon (;). For example:

```
Targets=HYPERV1;HYPERV2;
```

You can also edit the value for `CollectionIntervalInSeconds` based on your requirement. The default value is set to 300 seconds. For example:

```
CollectionIntervalInSeconds=300
```

Save and close the file.

8. Restart the HPE Cloud Optimizer Collector Service either by service management console or run the following command from command prompt:
`SC start HPE Cloud Optimizer Collector Service`
9. To uninstall the installer, run the following command:
`Uninstall.bat`

Note: While monitoring a Hyper-V environment, in case the Dynamic Memory is not enabled for VMs, the HPE Cloud Optimizer shows the memory utilization as 100%.

Proxy and Remote WMI Configuration

To provide WMI access on Proxy Side (Win2k12), follow these steps:

1. Go to **Server Manager > Local Server > Tools > Computer Management > Local User and Groups**.
2. Add the domain user to **Administrators** and **WinRemoteWMIUsers_** groups.

For information about providing DCOM remote access permissions, see [Set Computer-wide Launch and Activation Permissions](#).

Note: Domain user must be added to domain admin or local admin privileges due to Windows permissions. HPE Cloud Optimizer Collector Service does not make any changes to the domain or Hyper-V environment. It makes read-only requests for the system information.

Adding KVM Datasources

Prerequisites

- Install **Libvirt** and all its dependencies on the HPE Cloud Optimizer server.
Libvirt is an open source management tool to manage virtualized platforms such as Linux, KVM, XEN, and so on. The following table lists the versions of the different Linux types on which you can install Libvirt.

For KVM:

Linux Types	OS Versions
CentOS	6.4 (64-bit)
	6.5 (64-bit)
RHEL	6.4 (64-bit)
	6.5 (64-bit)
Ubuntu	12.04

- Install **Expect** and all its dependencies on the HPE Cloud Optimizer server. Expect is a tool that communicates with interactive programs, such as, `ssh`.

- The 22 (SSH) port should be open to allow HPE Cloud Optimizer collect data.

Note: To collect memory metrics, you need to install a lightweight collector on the KVM Hypervisor.

To add a KVM datasource, follow these steps:

1. Click **Collection & Analysis** option in the Settings page.
2. Click the **Add Datasource** link.
3. Select **KVM** option from the Domain drop-down.

If the KVM Hypervisor is not HPE Helion KVM Compute node, clear the Helion Compute Node check box. Here, HPE Helion KVM Compute node refers to KVM Compute node deployed from within HPE Helion OpenStack.

If the KVM Hypervisor is a HPE Helion KVM Compute node, ensure that the following pre-requisites are met:

- **Enable passwordless authentication between HPE Cloud Optimizer and HPE Helion KVM compute node**

Perform the following steps to enable passwordless authentication between HPE Cloud Optimizer and HPE Helion Compute node:

- A. On the HPE Cloud Optimizer machine, go to the following location:

`/root/.ssh`

- B. Open the `id-rsa.pub` file and copy the ssh public key.

- C. Log on to the HPE Helion KVM Compute node via the deployer node.

- D. Go to the following location:

`/home/stack/.ssh/authorized_keys`

- E. Open the `authorized_keys` file and append the HPE Cloud Optimizer public key to it.

- F. Save and close the file.

- G. Log on the HPE Cloud Optimizer machine.

- H. Run the following command to log on to the HPE Helion KVM Compute node from HPE Cloud Optimizer:

```
ssh stack@<compute_node>
```

In this instance,

`<compute_node>` is host name of the HPE Helion KVM Compute node.

You can now log on to the HPE Helion KVM Compute node from HPE Cloud Optimizer without authentication.

- **Run the following command on the HPE Helion KVM compute node:**

```
sudo apt install netcat-openbsd
```
 - **Download PVcAdvisor.**
 Click **Download PVcAdvisor** to download the HPPVcAdvisor-0.10.001-Linux2.6_64.rpm.gz file. Extract the file and install this rpm on KVM Hypervisor.
4. Type the IP address or host name of the datasource in the **IP/URL/Host Name** field.
 5. Type the user name corresponding to the specified datasource in the **User Name** field.
 6. Click **Test Connection**, if you want to validate the credentials provided.
 The system validates the credentials and displays a message stating whether the credentials are valid or not. The system also validates if the connection is successful or not.
 7. Click **Add**.
 The Status of Datasources list displays the added datasource.

Adding XEN Datasources

Prerequisites

- Install **Libvirt** and all its dependencies on the HPE Cloud Optimizer server.
 Libvirt is an open source management tool to manage virtualized platforms such as Linux, KVM, XEN, and so on. The following table lists the versions of the different Linux types on which you can install Libvirt.

For XEN:

Linux Types	OS Versions
Ubuntu	12.04
SLES	11 Service Pack 2, 11 Service Pack 3

- Install **Expect** and all its dependencies on the HPE Cloud Optimizer server. Expect is a tool that communicates with interactive programs, such as, `ssh`.
- The 22 (SSH) port should be open to allow HPE Cloud Optimizer collect data.

Note: The CPU and memory statistics are not available for XEN host remotely. You can view them on the **Dom0** guest instance.

To add a XEN datasource, follow these steps:

1. Click **Collection & Analysis** option in the Settings page.
2. Click the **Add Datasource** link.
3. Select **XEN** option from the Domain drop-down.
4. Type the IP address or host name of the datasource in the **IP/URL/Host Name** field.
5. Type the user name corresponding to the specified datasource in the **User Name** field.
6. Click **Test Connection**, if you want to validate the credentials provided.
The system validates the credentials and displays a message stating whether the credentials are valid or not. The system also validates if the connection is successful or not.
7. Click **Add**.

The Status of Datasources list displays the added datasource.

Note: The CPU, memory, and network statistics are not available for XEN host remotely.

Adding OpenStack Datasources

Prerequisites

- Before adding OpenStack datasources to HPE Cloud Optimizer, you must add the Hypervisors of OpenStack instances.
- We use Admin API endpoints of the Keystone and Nova components of OpenStack, so HPE Cloud Optimizer should be a part of that network, which has access to these API endpoints.

Adding Hypervisor Details to HPE Cloud Optimizer

To add a KVM or vSphere Hypervisor, follow these steps:

1. Click **Collection & Analysis** option in the Settings page.
2. Click the **Add Datasource** link.
3. Select **vCenter** or **KVM** option from the Domain drop-down.

If the KVM Hypervisor is not HPE Helion KVM Compute node, clear the **Helion Compute Node** check box. Here, HPE Helion KVM Compute node refers to KVM Compute node deployed from within HPE Helion OpenStack. To view the pre-requisites for adding a KVM datasource, see [Adding KVM Data Source](#).

If the KVM Hypervisor is a HPE Helion KVM Compute node, ensure that the following pre-requisites are met:

- **Enable passwordless authentication between HPE Cloud Optimizer and HPE Helion KVM compute node**

Perform the following steps to enable passwordless authentication between HPE Cloud Optimizer and HPE Helion Compute node:

- On the HPE Cloud Optimizer machine, go to the following location:
`/root/.ssh`
- Open the `id-rsa.pub` file and copy the ssh public key.
- Log on to the HPE Helion KVM Compute node via the deployer node.
- Go to the following location:
`/home/stack/.ssh/authorized_keys`
- Open the `authorized_keys` file and append the HPE Cloud Optimizer public key to it.
- Save and close the file.
- Log on the HPE Cloud Optimizer machine.
- Run the following command to log on to the HPE Helion KVM Compute node from HPE Cloud Optimizer:

```
ssh stack@<compute_node>
```

In this instance,

`<compute_node>` is host name of the HPE Helion KVM Compute node.

You can now log on to the HPE Helion KVM Compute node from HPE Cloud Optimizer without authentication.

- **Run the following command on the HPE Helion KVM compute node:**

```
sudo apt install netcat-openbsd
```

- **Download PVcAdvisor.**

Transfer and install it on the HPE Helion KVM compute node.


4. Type the IP address or host name of the hypervisor in the **IP/URL/Host Name** field.

Note: The KVM Hypervisor node of OpenStack must be added with the name that is displayed in the OpenStack horizon dashboard. You can get the name of the node in the following ways:

- Using the OpenStack horizon dashboard User Interface
- Using the Nova command line client

Using the OpenStack horizon dashboard User Interface

- a. In the OpenStack horizon dashboard user interface, go to **Admin > Hypervisors**.
- b. From the **Hypervisors** table, note the Hostname of the KVM Hypervisor node to be added.



The screenshot shows the OpenStack Horizon dashboard. The left sidebar has a red box around the 'Hypervisors' link. The main content area is titled 'All Hypervisors' and includes a 'Hypervisor Summary' section with three gauges: VCPU Usage (Used 6 of 34), Memory Usage (Used 5GB of 392GB), and Disk Usage (Used 6.0GB of 8.4TB). Below this is a table of hypervisors.

Hostname	Type	VCPU's (total)	VCPU's (used)	RAM (total)
domain-c15(Shire Cluster)	VMware vCenter Server	24	0	373GB
node014	QEMU	4	0	7GB
node043	QEMU	4	4	7GB

Using the Nova command line client

Type the following command in the Nova command line interface:

```
nova hypervisor -list
```

To run this command, you should have the Nova command line tool installed. The command returns the hypervisor hostnames.

Perform the following steps on the HPE Cloud Optimizer machine:

- a. Go to **/etc/hosts**.
- b. Add a IP address and hostname entry in the **/etc/hosts** file.

For example,

16.184.88.92	node014
16.184.76.43	node043

5. Type the user name of the hypervisor in the **User Name** field.
6. Type the password of the hypervisor in the **Password** field.
7. Click **Test Connection**, if you want to validate the credentials provided.

The system validates the credentials and displays a message stating whether the credentials are valid or not. The system also validates if the connection is successful or not.

8. Click **Add**.

The Status of Datasources displays the added KVM or vSphere Hypervisor.

After adding the hypervisor information, you can add the OpenStack datasource to HPE Cloud Optimizer.

Also, for HPE Cloud Optimizer to provide all the required data to an OpenStack user, the user should have admin rights for those tenants for which the user wants monitoring. Perform the following steps to give the user Admin rights:

Example:

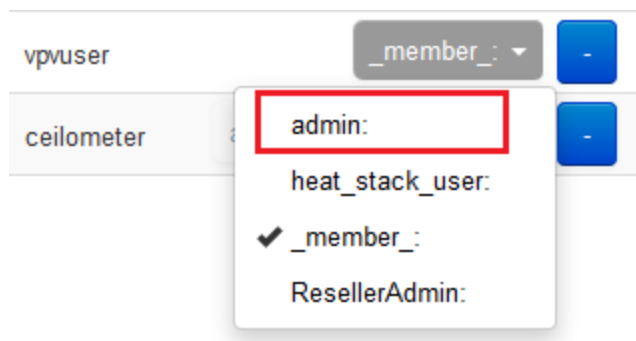
1. Log on to the OpenStack horizon dashboard as a HPE Cloud Optimizer user.
2. The Overview page displays three tenants as follows:

Project Name	VCPU's	Disk	RAM	VCPU Hours	Disk GB Hours
admin	3	41	4GB	34.69	474.04
Human_Resource	1	1	512MB	11.56	11.56
VPV_1	2	2	1GB	23.12	23.12
Displaying 3 items					

3. Log on to the OpenStack horizon dashboard as a Admin.
4. Go to **Identity Panel > Projects**.
5. To modify users for service tenant, in Actions column, click **Modify Users**.

<input type="checkbox"/>	Name	Description	Project ID	Enabled	Actions
<input checked="" type="checkbox"/>	Human_Resource	Human_Resource TestUsage	20c49071897b4cab86c1f9c8f13c883	True	Modify Users More *
<input type="checkbox"/>	demo	Demo Tenant	2526828be1c4a8a/ba62ed5a0b79054	True	Modify Users More *
<input type="checkbox"/>	VPV_1	VPV	5a3cd8d855c48708500c753f34a8a2a	True	Modify Users More *
<input type="checkbox"/>	admin	Admin Tenant	62764902c15440fcbac31bc3aee285	True	Modify Users More *

6. The Edit Projects dialog box appears. To give a member Admin access, click the project member and from the drop-down select **admin** and click **Save**.



7. This member is now an Admin for that tenant.

To add an OpenStack datasource, follow these steps:

1. Click **Collection & Analysis** option in the Settings page.
2. Select **OpenStack** option from the Domain drop-down.
3. Type the IP address or host name of the target datasource in the **Target IP/Host Name** field. This should be the IP address or host name from external API endpoint.
4. The **Port** field displays the default port number for OpenStack.
5. Type the admin user name corresponding to the specified datasource in the **User Name** field.
6. Type the admin password corresponding to the specified user name in the **Password** field.
7. Type the authentication token that you receive after the installing the OpenStack on the host in the admin **Auth Token** field.

To locate the authentication token:

- a. On the OpenStack system, go to **/etc/keystone**.
 - b. Open the **keystone.conf** file.
 - c. The authentication token is located in the **admin_token** field in the **keystone.conf** file.
8. Click **Test Connection**, if you want to validate the credentials provided.
The system validates the credentials and displays a message stating whether the credentials are valid or not. The system also validates if the connection is successful or not.

9. Click **Add**.

The Status of Datasources list displays the added datasource.

Note: If a datasource is down, HPE Cloud Optimizer does not collect data for that datasource.

Adding OneView Datasource

Prerequisites

The 5671 port should be open to allow HPE Cloud Optimizer to collect data from HPE OneView targets.

To integrate HPE Cloud Optimizer with HPE OneView, follow these steps:

On HPE Cloud Optimizer machine:

1. Log on as an administrator.
2. From the HPE Cloud Optimizer home page, select **Settings**.
3. Click **Add Datasource** option in the Settings page.
4. Select **OneView** from the **Domain** drop-down list.
5. Type the IP address or host name of the HPE OneView appliance in the **OneView IP/Host Name** field.
6. Type the user name of the HPE OneView appliance in the **User Name** field.

Note: The user to be configured in HPE Cloud Optimizer should have PUT and POST access to the following URL:

```
https://<IP address of HPE OneView appliance>/rest/login-sessions
```

7. Type the password of the HPE OneView appliance in the **Password** field.
8. Click **Test Connection**, if you want to validate the credentials provided. The system validates the credentials and displays a message stating whether the credentials are valid or not. The system also validates if the connection is successful or not.
9. Click **Add**. The Connections list displays the added HPE OneView appliance.

Note: If there is any change in HPE OneView setup, such as addition or deletion of Enclosures and Blade Servers, HPE Cloud Optimizer recommends restarting the data collection for the specific HPE OneView setup. This ensures that the Treemap is updated with the latest changes in HPE OneView setup.

Adding Amazon Web Services Datasource

HPE Cloud Optimizer includes support for Amazon Web Services (AWS) Elastic Compute Cloud (EC2) performance monitoring. Amazon Web Services or AWS is the cloud computing platform solutions from Amazon. By adding support to manage the performance and capacity of AWS EC2, HPE Cloud Optimizer moves closer to its vision of a managing multiple hypervisors and competing clouds. With AWS monitoring, you

can ensure the reliability, availability, and performance of your Amazon EC2 instances and AWS solutions.

Note: HPE Cloud Optimizer's Forecast feature is not available with AWS Previous Generation Template.

Prerequisites

To configure web proxy on the HPE Cloud Optimizer machine in a private network:

1. Log on as a root user.

2. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

In this instance, `<install_dir>` is the directory where HPE Cloud Optimizer is installed. The default location is `/opt/OV`.

A text file opens.

3. In the text file, create the following namespace:

```
[bbc.http]
```

In the `[bbc.http]` namespace, type the following:

```
[bbc.http]
```

```
PROXY=<Proxy_Name>:<Port_Number>+(cloudwatch.aws.com)
```

In this instance,

`<Proxy_Name>` is the name of the web proxy.

`<Port_Number>` is the port number of the web proxy.

`cloudwatch.aws.com` -This is a constant string value that is required by HPE Cloud Optimizer to connect to AWS instances.

If you want to enable or disable proxy for specific IP addresses, use the following command:

```
PROXY=<Proxy_Name>:<Port_Number>+(cloudwatch.aws.com,<IP address of Amazon instances>)
```

In this instance,

`<IP address of Amazon instances>` is the IP address of the Amazon instance. You can give multiple IP addresses separated by a comma. You can specify IP addresses as explained in the following example:

```
PROXY=web-proxy:8088-(*.hp.com)+(*.bbn.hp.com;*),
```

The proxy, `web-proxy` is used with port 8088 for every server (*) except hosts that match `*.hp.com`, for example `www.hp.com`. If the hostname matches `*.bbn.hp.com`, for example `merlin.bbn.hp.com` the proxy server is used. It is also possible to use IP addresses instead of host names, so `15.*.*.*` or `15:*.***:*.***` is valid as well, however, the correct number of dots or colons must be specified.

To add an AWS datasource, follow these steps:

1. In the **Settings** page, click **Collection & Analysis**.
2. Click **Add Datasource** link.
3. Select **AWS** option from the Domain drop-down.
4. Type account ID, user name, access key, and secret key in the **Account ID**, **User Name**, **Access Key**, and **Secret Key** fields.

Click **Clear** to clear any input.

Note:

Account ID - The AWS account ID for the user.

User Name - The AWS User Name for the user.

Access Key - The AWS assigned Access Key for a secure user request.

Secret Key - The AWS assigned Secret Key for a secure user request.

5. Click **Test Connection**, if you want to validate the credentials provided.
6. Click **Add**.

The Status of Datasources list displays the added datasource.

Adding Physical Server as a Datasource

In addition to monitoring your virtual infrastructure, HPE Cloud Optimizer has been enhanced to monitor non-virtualized or physical servers as well. With the capability of monitoring physical servers, HPE Cloud Optimizer now gives comprehensive top-to-bottom visibility and control of your infrastructure.

With network environments becoming increasingly diverse and complex, there is a continuous need to look at opportunities to convert physical infrastructure into virtual. With HPE Cloud Optimizer, you can analyze and forecast the capacity for various physical resources being monitored. The accurate forecast information for physical servers enables you to plan the migration from physical to virtual infrastructure efficiently.

Using physical server monitoring in HPE Cloud Optimizer, you can do the following:

- View the topology of your physical servers with the Treemap view. Use the Treemap to identify and view the problems at a glance.
- Triage analysis with the Workbench and gain the ability to drill down for a rich set of metrics.
- Analyze and forecast future capacity needs.
- View the list of servers, which are running and which are not running.

Prerequisites

The 8081 (HTTP) port should be open to allow HPE Cloud Optimizer collect data.

To add a Physical Server as a Datasource, follow these steps:

1. Click **Collection & Analysis** option in the Settings page.
2. Click the **Add Datasource** link.
3. Select **Physical Server** from the Domain drop-down list.
4. Click **Download Collector for Windows** or **Download Collector for Linux** link to download the respective tar files:

On Windows: `vPVPhysMon_windows.tar`

On UNIX: `vPVPhysMon_linux.tar`

Note: Installation of the Physical Server collector is not supported on the machine where Real Time Guest OS Drill Down is already configured.

5. Copy the downloaded tar file to the physical server that is to be monitored and then extract the contents. An `hpcs` folder containing the integration packages is created. Copy the `hpcs` folder to the respective location:

On Windows: `c:\temp`

On UNIX: `/tmp`

6. Install the collector. To install, go to the following location:

On Windows: `c:\temp\hpcs\`

On UNIX: `/tmp/hpcs`

7. Run the following commands on the Physical Server:

- a. To install the packages:

- On Windows: `cscript setuphpcs.vbs -install`

- On UNIX: `./setuphpcs.sh -install`

- b. To remove the packages:

- On Windows:

To remove the package from a Windows node, follow these steps:

A. Click **Start**. Go to **Control Panel > Programs > Programs and Features > Uninstall a program**.

B. Right-click the program **HPE ComputeSensor** and then click **Uninstall**. If you are prompted to confirm the removal of the program, click **Yes**.

- On Linux: `rpm -e HPComputeSensor`

- On Linux Debian: `dpkg -P hpcomputesensor`

8. After the installation is complete, HPE Cloud Optimizer starts the data collection for

the physical server. You can see the status of data collection under **Settings> Collection & Analysis** option. To see the status of data collection, click the **PhysicalServers** link in the table. **Physical Server Collection Status** page appears.

The **Physical Server Collection Status** page has the **Running** and **Stopped** tabs. In the **Running** tab, you can see the list of physical servers from which data is being collected. In the **Stopped** tab, you can see the list of physical servers from which data is no longer collected.

Verifying the Installation of Physical Server Collector

To verify if the installation of Physical Server Collector is successful, check for the following:


1. Go to the following directory:
On Windows nodes: %HPCSInstallDir%
On Linux or Linux-Debian nodes: /opt/OV/hpcs
2. Check the **hpcstrace.log** file. The following message appears:
Started successfully. Type the URI to connect --> http://localhost:381

Note: On Windows nodes: A new entry named HPE Compute Sensor Service will be listed under **Services.msc**.

Monitoring Physical Servers using HPE Cloud Optimizer Features

You can monitor the physical servers in HPE Cloud Optimizer using the following features:

Feature	Description
Treemap	<p>When a physical server is added to HPE Cloud Optimizer, you can monitor the performance of the server in Treemap.</p> <p>Select Physical Server as the domain in the Domain list of Treemap. Treemap appears with the performance information about the physical servers added to HPE Cloud Optimizer.</p>

Feature	Description
	 <p>When Physical Server is selected as the domain, you can perform the following tasks from the Treemap:</p> <ul style="list-style-type: none"> • View the number of physical servers being monitored in the Resource pane • Access Forecast to view CPU and Memory usage forecast for physical servers • Access Workbench to view the performance metrics for physical servers

Feature	Description
Workbench	<p>To view the performance graphs for physical servers, you can launch the Workbench feature from the physical server in the Treemap.</p> <p>In the Treemap, right-click a server and select Launch Workbench. The Workbench page appears with the selected server name in the Configuration Items list.</p> <p>You can perform the following tasks from the Workbench:</p> <ul style="list-style-type: none"> • View the performance metrics for the physical server • Visualize and triage the performance of the physical servers • Draw and view graphs with the help of available metrics. <p>Additionally, you can save frequently used graphs as favorites.</p> <p>For more information on Workbench, see the Performance Graphs section.</p>
Forecast	<p>To view the forecast for physical servers, you can launch the Forecast feature from Treemap.</p> <p>In the Treemap, right-click a physical server and select Forecast. The Forecast page appears for the selected physical server.</p> <p>You can view the forecast for CPU usage and Memory usage of the physical server.</p> <p>For more information on Forecast, see the Forecast section.</p>

Monitoring the Health of a Datasource

To view the health of a VMware datasource, click the IP in the IP/URL/Host Name header of the Status of Datasources table. The Datasources Health page appears.

- Viewing Collection Information

The Collection Information section provides you information about the data retention by HPE Cloud Optimizer. The section also provides you detailed information about the duration for which the data is retained in the database. The data collected is summarized, and rollup is done to show the consolidated data. The data retention period varies, based on the license applied. For more information on the different license types, see the Installation Guide.

To set the Collection Interval, select **5 minutes** or **15 minutes** options.

- Viewing Capacity Analysis

The Capacity Analysis section provides you information on when the last analysis was run and when the next analysis is planned.

To start Capacity Analysis, select the **All time** or **Last month** or **Last week** option, and then click **Run Analysis**.

Restarting Collection

To restart the datasource collection:

1. Click **Collection & Analysis** option in the Settings page.
2. Select the datasource that you want to restart.
3. Click **Restart** from Actions.

Note: To restart Hyper-V, run the following command on the proxy system:

```
restart HPE Cloud Optimizer collector service
```

License

The License option provides information about your HPE Cloud Optimizer licenses. You can also use this option to start using your Evaluation License as well as import a HPE Cloud Optimizer license. To purchase the respective license, click **Buy License** link. Enter your details in the **Contact Me** tab and click **Submit**. Make sure email client is configured in your browser.

For more information on the difference between the licenses and the features available, see *HPE Cloud Optimizer Installation Guide*.

The following table lists the sections available in the License option.

Section	Description
License Status	<p>The License Status section lists the following information about your current HPE Cloud Optimizer license:</p> <ul style="list-style-type: none">• Maximum number of OS instances that are provided by the active licenses• Number of OS instances that are currently being monitored by the active licenses• Data Retention• If you have a Server Instance License enabled, you can see the total

Section	Description
	<p>Server Instance capacity, the current Server Instance Count, and the OS instances that are covered by the Server license.</p> <ul style="list-style-type: none"> License Alert Threshold - The License Alert Threshold slider allows you to set the threshold percentage for the licensed OS instance capacity. When the current OS instances reach the set threshold percentage, HPE Cloud Optimizer displays alert messages with the current usage of the OS instance capacity and the remaining OS instance count. This helps you to start evaluating license purchase decisions well in time. You can set this value between 60% - 100%. <p>For example, if your total OS instance capacity is 2000 and License Alert Threshold is set as 90%, HPE Cloud Optimizer starts displaying alert messages when the current OS instance count reaches 1800.</p>
Installed License List	<p>The Installed Licenses List section lists the following information regarding the active and all installed licenses for HPE Cloud Optimizer:</p> <ul style="list-style-type: none"> Type of license installed Date of license expiry

Note: The number of OS instances denote only the VMs and hosts.

Manage License

To import the license, perform the following tasks:

1. Go to the **License** option in Settings page.
2. Click **Import License** link. The Manage License page appears.
3. Enter the license key and click **Import License**.

To clear the license key typed in, click **Reset**.

Monitoring License

To start monitoring only license, click the **Switch to 'Monitoring-only' Mode** link.

If you want to use HPE Cloud Optimizer's monitoring capabilities only, you can enable the Monitoring License. With the Monitoring license, you can only use limited set of features in the **Settings > License**. All the other HPE Cloud Optimizer features are disabled.

You can enable the Monitoring License from the user interface and from the command line.

From User Interface:

To enable Monitoring License from the user interface, follow these steps:

1. Go to the **License** option in the **Settings** page.
2. Click the **Switch to 'Monitoring-only' Mode** link.

From Command Line:

To enable Monitoring License from the command line, follow these steps:

1. On the HPE Cloud Optimizer machine, log on as a root user.
2. Run the following command to activate the Monitoring Only license:

```
pv startmonitoringonly
```

The Evaluation, Premium, and Express licenses override the existing Monitoring license. After the end of respective license entitlement periods, HPE Cloud Optimizer reverts to the Monitoring license features. Similarly, if you are currently using HPE Cloud Optimizer in evaluation mode, and you enable the Monitoring License, all the features would be available only for the evaluation period. After the Evaluation license validity expires, the Monitoring license continues to be active.

Integrating HPE Cloud Optimizer with Other HPE Products

You can use the Integration option to integrate HPE Cloud Optimizer and use it with products such as HPE Performance Manager (PM), HPE Business Service Management (BSM), HPE Cloud Service Automation (CSA), HPE Operations Manager (HPOM), and Real Time Guest OS Drill Down.

Note: The Integration feature is available only with the Evaluation and Permanent licensed versions of HPE Cloud Optimizer.

Navigation

To access the feature, Click **Settings > Integration**.

The following section provides information about the integration of HPE Cloud Optimizer with other HPE Software Products:

Use	Integration
Consolidates the events using HPE Operations Bridge (HPE OpsBridge) and displays it on the HPE Cloud	Integrating HPE Cloud Optimizer with BSM

Use	Integration
Optimizer interface.	
Monitor the performance and analyze the capacity, usage, and forecast trends of the virtualized infrastructure that are custom views for the CSA user.	Integrating HPE Cloud Optimizer with CSA
View the Alert messages on HPOM console to take quick actions.	Integrating HPE Cloud Optimizer with HPOM
Drill down to check the health of the VMs. Get the process details which are consuming more resources (CPU or memory).	Integrating HPE Cloud Optimizer with Real Time OS Drill Down
Monitor the performance and analyze the capacity, usage, and forecast trends of HPE OneView Blades and Enclosures.	"Integrating HPE Cloud Optimizer with HPE OneView" on page 55

Integrating HPE Cloud Optimizer with HPE Cloud Service Automation (HPE CSA)

Integrating HPE Cloud Optimizer with HPE CSA enables the CSA user to monitor the performance and analyze the capacity, usage, and forecast trends of the virtualized infrastructure that are custom views for the CSA user. As an administrator, you must add the credentials of the CSA user from HPE Cloud Optimizer administrator tab. The CSA user (operator) can launch the custom views after logging into the system.

As a CSA user, you can perform the following tasks on the available custom views:

- Monitor the aggregated resource utilization of the available VMs
- Overview about resource utilization trend of all the VMs
- Get recommendations to fine tune resource allocation (CPU, memory) for optimal virtual machine performance
- Drill down to see forecast for each virtual machine for CPU, memory, and storage utilization trend for 90 days

As an administrator, complete the following tasks in order:

Tasks	Description
Check the performance of the VMs by logging on as CSA user.	Integrating HPE Cloud Optimizer with CSA from HPE Cloud Optimizer Console.

Tasks	Description
Access to CSA user to avoid multiple log on.	Configuring CSA for HPE Cloud Optimizer Single Sign-On Utility.

After completing these tasks, you can log on as a CSA user. For more information, see [Monitoring the Infrastructure as a Cloud Service Automation \(CSA\) user](#).

Integrating HPE Cloud Optimizer with CSA from HPE Cloud Optimizer Console

To integrate HPE Cloud Optimizer with CSA, follow the steps:

1. Log on to HPE Cloud Optimizer as an administrator.
2. Add the datasources to HPE Cloud Optimizer that are used as a provider by CSA to deploy the virtual machines.
3. Enable LDAP. For more information, see [Integrating with Microsoft Active Directory and OpenLDAP](#).
4. Type the URL for which you want to integrate CSA with HPE Cloud Optimizer in the **CSA URL** field.
For example: `https://<CSA Hostname or IP address>:8444`
5. Type the administrator user name and password in the **User Name** and **Password** fields respectively.
6. Clear the **Deploy Business Group** check box if you do not want to deploy **CSA** Business Grouping. If the **Deploy Business Group** option is selected, a default Business Grouping named **CSA** is deployed in HPE Cloud Optimizer. The **Deploy Business Group** option is selected by default.
7. In the **Server Component Types** box, type the name of the custom server component type. Select the **Exact Match** check-box if you know the exact value for Server Component Types.
8. Instance Mapping Rule section allows you to define the rule, which you want the VMs to reconcile with HPE Cloud Optimizer. Here the property name, which will contain the value to be used to match with the data of the HPE Cloud Optimizer is required. The default value is the `hostName, Name`.

In the **CSA Component Property** box, type the property name of the CSA component. This property name contains values, which will reconcile with HPE Cloud Optimizer. The value of this property will be matched with the HPE Cloud Optimizer **Cloud Optimizer Server Property** value.

Instance Mapping Rule

CSA Component Property : ServerID ⓘ

Cloud Optimizer Server Property : System ID ⓘ

Add

CSA Component Property	Cloud Optimizer Server Property	Remove
hostName	Name	
ServerID	System ID	✕

Note: For OpenStack or Helion, HPE Cloud Optimizer reconciles the VM's UUID. So, ensure that VM's property, which contains the UUID is mentioned in the **CSA Component Property** box.

- From the **Cloud Optimizer Server Property**, select the appropriate attribute.
Name - This is the display name of the VM in the Hypervisor.
IP Address - This is the IP address of the VM.
System ID - This is the UUID of the VM.
- Click **Save**. A message stating that the CSA information is saved successfully appears.

Note: If you change any configuration property and would like to save the CSA information again, it is highly recommended to reset CSA users information using the **Reset CSA Users Info** button before saving.

After the configuration, CSA is integrated with HPE Cloud Optimizer. The CSA user (s) can check the performance of their VMs by logging on as CSA user.

Monitoring the Infrastructure as a Cloud Service Automation (CSA) user

As an administrator, you can complete the configuration steps to create a CSA user. For more information, see the steps in ["Integrating HPE Cloud Optimizer with Other HPE Products" on page 39](#).

Follow the steps to log on as CSA user:

- Browse to the HPE Cloud Optimizer login page.
- Type the CSA user name in the **Username** field.
- Type the password in the **Password** field.
- Select **CSA** from the **Context** field.
- Select the required organization from the **Organization** field.

Note: HPE Cloud Optimizer supports only those CSA organizations that have configured authentication using LDAP. For example, HPE Cloud Optimizer will not support the out-of-the-box CSA organizations such as **Consumer** and **Provider**, if they have not set the LDAP authentication.

6. Click **Login**.

When a CSA user logs on to HPE Cloud Optimizer, the user needs to select the context as CSA and also the corresponding organization. HPE Cloud Optimizer would then launch a context-sensitive view of capacity usage and forecast trend details for infrastructure owned by the CSA consumer.

Configuring CSA for HPE Cloud Optimizer Single Sign-On Utility

As a CSA user you can configure HPE Cloud Optimizer for Single Sign-On (SSO) thereby eliminating the process of multiple sign-in.

Follow these steps:

1. Log on as a root user.
2. Go to `lwssofmconf.xml` file at the following location:
`/opt/OV/www/webapps/PV/WEB-INF/classes/`
3. Edit the values of `<domain>` and `<initString>` in the `lwssofmconf.xml` file.

For example:

```
<validation>
<in-ui-lwsso>
<lwssoValidation id="ID000001">
<domain><hp.com></domain>
<crypto cipherType="symmetricBlockCipher"
engineName="AES" paddingModeName="CBC" keySize="256"
encodingMode="Base64Url"
initString=<mention the init string here></crypto>
</lwssoValidation>
</in-ui-lwsso>
</validation>
```

Make sure the `<domain>` and `<initString>` is the same as mentioned in the HPE SSO configuration.

4. Run the following command to restart tomcat:
`/opt/OV/bin/ovc -retstart ovtomcatB`
5. **Optional:** Login to CSA user page and give the following path in a new tab:
`http://<FQDN>:8081/PV/?CTX=CSA&CSA_ORG=<ID_Of_The_Organisation>`
Here, `<FQDN>` is the domain mentioned in `/opt/OV/www/webapps/PV/WEB-INF/classes/lwssofmconf.xml` file and `<ID of the organization>` is the ID of the organization and not the name of the organization.

Integrating BSM with HPE Cloud Optimizer

To integrate BSM with HPE Cloud Optimizer, perform the following steps:

On the BSM Machine:

Obtain the **BSM Token Creation Key** from the BSM user interface. The key can be accessed from **Admin > Platform > Users and Permissions > Authentication Management** from the BSM user interface.

On the HPE Cloud Optimizer Machine:

1. Log on as a root user.
2. Go to `lwssofmconf.xml` file at the following location:
`/opt/OV/www/webapps/PV/WEB-INF/classes/`
3. Edit the values of `<domain>` and `<initString>` in the `lwssofmconf.xml` file.

For example:

```
<validation>
<in-ui-lwss>
<lwssValidation id="ID000001">
<domain><hp.com></domain>
<crypto cipherType="symmetricBlockCipher"
engineName="AES" paddingModeName="CBC" keySize="256"
encodingMode="Base64Url"
```

```
initString=<mention the init string obtained from the BSM
Machine></crypto>
```

```
</lwsssoValidation>
```

```
</in-ui-lwssso>
```

```
</validation>
```

Make sure the *<domain>* and *<initString>* are the same as mentioned in the HPE SSO configuration.

4. Use the TRUSTED_SITE_URL parameter to configure the trusted website for HPE Cloud Optimizer. An application will be able to embed HPE Cloud Optimizer page within itself only if the *TRUSTED_SITE_URL* is set in the *PVconfig.ini* file.

For example, if the application that wants to embed HPE Cloud Optimizer has the URL *mydashboard.company.com/home/*, you must do the following configuration in the *PVconfig.ini* file:

```
TRUSTED_SITE_URL=https://<application_URL, for example,
mydashboard.company.com>
```

5. Restart ovtomcatB using the command:

```
/opt/OV/bin/ovc -restart ovtomcatB
```

On the BSM machine, use the following URL to include the HPE Cloud Optimizer component.

```
http://<vPV_Hostname or IP address>:<port number>/PV/?CTX=BSM
```

For more information on including components in BSM, refer *HPE Business Service Management* documentation.

Integrating HPE Cloud Optimizer with HPE Operations Manager

You can integrate HPE Cloud Optimizer with HPOM to view the alerts on HPOM console.

Follow the steps for integration:

Note: You can start the integration of HPE Cloud Optimizer and HPOM only after you have completed the installation of HPE Cloud Optimizer on a node.

On HPE Cloud Optimizer machine:

1. Log on as an administrator.
2. From the HPE Cloud Optimizer home page, select **Settings > Integrations**.
3. Download the following packages are available on the location:
 - Windows - **VPV-OMW-Integration.zip**

- UNIX - **VPV-OMx-Integration.tar**

Note: Make sure to transfer the integration packages to HPOM management server.

4. Go to ["On HPOM Management Server" below](#) and follow the instructions.
5. Install HPE Operations agent 11.13 or above on HPE Cloud Optimizer machine. You can install HPE Operations agent manually on the node or remotely from HPOM management server. For more information, see *HPE Operations agent Installation guide*.
6. Run the command to ensure the certificate is set between the management server and HPE Cloud Optimizer machine.
 - a. Run the command: `/opt/OV/bin/OpC/install/opcactivate -srv <hostname> -cert_srv <hostname>`
In this instance, <hostname> is the hostname of the management server.
 - b. Go to HPOM management server and grant the certificate.
7. Deploy the following policies from HPOM server:

Policy	Description
vPV-DiscoverNodes	Discovers the available nodes.
vPV-TopologyDiscovery	Discovers the available infrastructure.
vPV-EventMonitor	Monitors events or alarms from vCenter server.

8. After all the nodes are discovered, then only deploy the following policies:

vPV-OMIntegration	Forwards the alerts to HPOM server.
vPV-AlertSensitivity	Alert forwarding sensitivity. By default, the medium sensitivity is deployed. <div> <p>Note: You can disable sensitivity only from HPE Cloud Optimizer machine. For more information, see Sensitivity.</p> </div>

On HPOM Management Server

Follow the steps:

Note: Make sure that you have installed all the available patches and hotfixes for HPOM:

- HPE Operations Manager for UNIX - 09.10; Patch IDs - OML_00071 + OML_00072 + OML_00073
- HPE Operations Manager for Windows - 9.10; Patch IDs - OMW_00178

1. Extract the contents of **VPV-OMW-Integration.zip** or **VPV-OMX-Integration.tar** folder. Copy the Integration package to the respective folder:

On Windows: c:\temp

On UNIX: /tmp

Note: No special settings are required for HPOM for UNIX.

Make sure that you check the following settings for HPOM for Windows:

- RAM - 16 GB
- Java Heap - (Go to the location:
HKLM\SYSTEM\CurrentControlSet\services\OvAutoDiscovery
Server\JavaOptions and set the Xmx value to 8192)

2. Run the following commands to install or remove the HPE Cloud Optimizer integration packages:

a. Installing the packages:

- On Windows: `cscript install-integration-pack.vbs -install`
- On UNIX: `./install-integration-pack.sh -install`

b. Removing the packages:

- On Windows: `cscript install-integration-pack.vbs -remove`
- On UNIX: `./install-integration-pack.sh -remove`

Note: Before you remove the packages from the management server, make sure you remove the HPE Cloud Optimizer policies deployed on the HPE Cloud Optimizer nodes.

3. After the installation is complete, you can check the log file (install.log). This is available in the same location where the integration installation packages are available.

During the installation of the integration package on **HPOM for Solaris**, the following error appears:

```
"/etc/opt/OV/share/conf/OpC/mgmt_sv/integration/cfgupld/post/cvp_
upload.sh: [: not found"
```

You can ignore the error. The installation on HPOM for Solaris is successfully completed.

4. Add the HPE Cloud Optimizer machine to the **Nodes** folder on HPOM management server. For more details, see HPOM for Windows and UNIX online help.
5. Go to HPE Cloud Optimizer machine and perform [Step 5](#) to complete the steps.

Integrating HPE Cloud Optimizer with HPE Operations Manager *i* (OMi)

You can integrate HPE Cloud Optimizer with OMi to view the HPE Cloud Optimizer alerts, topology, and performance graphs on OMi. HPE Cloud Optimizer integration with OMi helps you to seamlessly monitor your virtual infrastructure from OMi.

Note: Integrating HPE Cloud Optimizer with OMi allows you to monitor the virtual infrastructure provided by VMware vCenter, Microsoft Hyper-V, and KVM.

This integration can be achieved by installing the **OMi Management Pack for vPV** on **HPE Business Service Management (HPE BSM) version 9.25 and HPE Operations Manager *i* (OMi) 10.01**. To support integration with HPE BSM 9.25, you must install the HPE BSM 9.25 hotfix. Contact support for more information.

Note: The **OMi Management Pack for vPV** should be installed on both the Gateway and the Data Processing server of OMi.

Prerequisites

All the monitored VM nodes should have the host name set and the VMware tools running on them.

The **OMi Management Pack for vPV** works with OMi and provides the following features:

- Ready to deploy Management Template to suit different monitoring requirements
- Ability to view HPE Cloud Optimizer alerts and topology on OMi
- Ability to view performance graphs

The **OMi Management Pack for vPV** includes the following artifacts:

Artifact	Name	Description
Management Template	vPV Monitor (300.0)	Monitoring the performance of virtual nodes. It consists of two Aspects: vPV Alerts and vPV Discovery.

Artifact	Name	Description
Aspects	vPV Alerts (300.0)	Monitors events or alarms from the vCenter server. It has four policies: vPV Alert Sensitivity, vPV Custom Alert Sensitivity Definition, vPV-EventMonitor, and vPV-OMIntegration.
	vPV Discovery (220.0)	Discovers the virtualization infrastructure.
Policies	vPV Alert Sensitivity	Used to handle alerts with high, medium, low, and custom sensitivity levels. The policy has a parameter named SensitivityLevel . SensitivityLevel can have values such as High, Medium, Low, and Custom.
	vPV Custom Alert Sensitivity Definition	If you want to use custom sensitivity, set the SensitivityLevel parameter as Custom and edit the vPV Custom Alert Sensitivity Definition policy.
	vPV-EventMonitor	Used to monitor events and alarms from the HPE Cloud Optimizer server.
	vPV-OMIntegration	Used to forward HPE Cloud Optimizer alerts to the HPE OM and OMi. You can use this policy to customize the severity of alerts. For more information on configuring the severity of alerts, see Customize the severity of Alerts .
	vPV-Discovery	Used to discover the virtualization infrastructure.
	vPV-SuppressAlerts	Used to suppress alerts.

Follow these steps for integration:

On HPE Cloud Optimizer machine:

1. Log on as an administrator.
2. From the HPE Cloud Optimizer home page, select **Settings > Integration**.
3. In the **HPE Operations Manager/HPE Operations Manager i (HPOM/HPE OMi)** section, click **Download Zip Package (OMi)**.
4. Copy the **vPV_OMi_Integration.zip** file and transfer it to the HPE OMi machine.
5. Install HPE Operations agent 11.14 on the HPE Cloud Optimizer machine. For more information, see the *HPE Operations agent Installation guide*.
6. Run the following command on the HPE Cloud Optimizer machine to ensure the certificate is set between HPE OMi and HPE Cloud Optimizer machine.
 - a. Run the command: `/opt/OV/bin/OpC/install/opcactivate -srv <hostname> -cert_srv <hostname>`
In this instance, <hostname> is the hostname of the management server.
 - b. Go to the HPE OMi machine and grant the certificate.

Note: The HPE OMi machine here refers to HPE OMi available with the HPE Business Service Management (HPE BSM) with Monitoring Automation version 9.25 (with BSM 9.25 hotfix applied) or OMi 10.01.

On the HPE OMi Machine:

Note: The following steps must be performed on both the Gateway and the Data Processing server of OMi.

Follow the steps:

Note: Make sure you have installed the **HPE OMi Management Pack for Infrastructure version 1.11** on the HPE OMi machine. You can download it from [HPE Live Network](#).

1. Extract the contents of the vPV OMi Integration Pack. Copy the Integration package to the respective folder:
On Windows: `c:\temp`
On UNIX: `/tmp`
2. Run the following commands to install the vPV OMi Integration Pack:
 - On Windows: `cscript vPV_integration_install.vbs -i`
 - On UNIX: `./vPV_integration_install.sh -i`

3. After the installation is complete, check the log file (install.log). This is available in <BSMInstallDir>/log/vpv_int_install.log.
4. Add the HPE Cloud Optimizer machine to the **Nodes with Operations Agent** group on the HPE OMi machine. To access,
On BSM, click **Admin > Operations Management > Setup > Monitored Nodes > Nodes with Operations Agent**.
On OMi, click **Administration > Setup and Maintenance > Monitored Nodes > Nodes with Operations Agent**.
5. Before you begin monitoring, deploy the Management Template or Aspects provided by the **OMi Management Pack for vPV**.

Note: If you are using **Monitoring Automation for Composite Applications** license, you can deploy either the **OMi Management Pack for vPV** Management Templates or **OMi Management Pack for vPV** Aspects to the CIs. For more information about deploying the HPE Cloud Optimizer OMi Integration Management Template, see [Identifying and Deploying HPE Cloud Optimizer OMi Integration Management Templates](#).

If you are using **Monitoring Automation for Server** license, you can deploy the **OMi Management Pack for vPV** Aspects. For more information about deploying the **OMi Management Pack for vPV** Aspects, see [Deploying the HPE Cloud Optimizer OMi Integration Aspects](#).

- **Deploying the Aspects**

To deploy the **OMi Management Pack for vPV** Aspects to the CIs, follow these steps:

Note: Skip this task if you are using **Monitoring Automation for Composite Applications** license and have assigned the **OMi Management Pack for vPV** Management Templates to the CI.

- i. Open the Management Templates and Aspects manager:
On OMi, click **Administration > Monitoring > Management Templates and Aspects**.
- ii. In the **Configuration Folders** pane, click **Configuration Folders > Infrastructure Management > vPV Management Pack**.
- iii. In the **Management Templates and Aspects** pane, right-click the **vPV Discovery** Aspect, and then click **Assign and Deploy** Item. The **Assign and Deploy** wizard opens.

Note: You must first deploy the **vPV Discovery** Aspect followed by the **vPV Alerts** aspect.

- iv. In the **Configuration Item** tab, click the CI to which you want to assign the Management Template, and then click **Next**.
- v. No editing is required in the **Required Parameters** tab. Click **Next** to go to **Required Parameters** tab.
- vi. To change the default values of the parameters, you can select the parameter and then click **Edit**. The **Edit Parameter** dialog box opens. Click **Value**, specify the value, and then click **OK**.
- vii. *Optional*. In the **Configure Options** tab, if you do not want to enable the assignment immediately, clear the **Enable Assigned Objects** check box. You can then enable the assignment later using the **Assignments & Tuning** manager.
- viii. Click **Finish**.

- **Identifying and Deploying Management Template**

To deploy the **OMi Management Pack for vPV** Management Templates to the CIs, follow these steps:

- i. Open the Management Templates and Aspects manager:
On OMi, click **Administration > Monitoring > Management Templates and Aspects**.
- ii. In the **Configuration Folders** pane, click **Configuration Folders > Infrastructure Management > vPV Management Pack**.
- iii. In the **Management Templates and Aspects** pane, right-click **vPV Monitor** Aspect, and then click **Assign and Deploy** item. The Assign and Deploy wizard opens.
- iv. In the **Configuration Item** tab, click the CI to which you want to assign the Management Template, and then click **Next**.
- v. No editing is required in the **Required Parameters** tab. Click **Next** to go to **Required Parameters** tab.
- vi. To change the default values of the parameters, you can select the parameter and then click **Edit**. The **Edit Parameter** dialog box opens. Click **Value**, specify the value, and then click **OK**.
- vii. *Optional*. In the **Configure Options** tab, if you do not want to enable the assignment immediately, clear the **Enable Assigned Objects** check box. You can then enable the assignment later using the **Assignments & Tuning** manager.
- viii. Click **Finish**.

Verifying the HPE Cloud Optimizer Integration with OMi

After the Management Template has been successfully deployed, you can view the HPE Cloud Optimizer alerts, topology, and performance graphs in OMi.

To verify the HPE Cloud Optimizer and OMi integration, do the following on the OMi machine:

- Verify if the **HPOprvPV** Package is available in the OMi Package Manager
- Verify if the **vPV_Infrastructure** view is created in the OMi Modeling Studio

To verify if the **HPOprvPV** Package is available in the OMi Package Manager, follow these steps:

1. Navigate to
On BSM, **Admin > RTSM Administration > CI Type Manager**.
On OMi, **Administration > RTSM Administration > CI Type Manager**.
2. Click the **Administrator** tab and select **Package Manager**.
3. Search for **HPOprvPV** Package in the list of available packages.

To verify if the **vPV_Infrastructure** view is created in the OMi Modeling Studio, follow these steps:

1. Navigate to
On BSM, **Admin > RTSM Administration > CI Type Manager**.
On OMi, **Administration > RTSM Administration > CI Type Manager**.
2. Click the **Modeling** tab and select **Modeling Studio**.
3. In the left pane, select **Resources**. Under the **Root** folder, locate **Operations Management**.
4. Under the **Operations Management** folder, look for **vPV > vPV_Infrastructure** view.

Viewing Performance Graphs in HPE OMi

When you integrate HPE Cloud Optimizer with HPE OMi, you can view performance graphs for VMs and hosts in HPE OMi.

Prerequisites

- Configure HPE OMi for HPE Cloud Optimizer Single Sign-On (SSO) Utility

Note: The HPE OMi machine here refers to HPE OMi available with the HPE Business Service Management (HPE BSM) with Monitoring Automation version 9.25 and HPE OMi 10.01. To support integration with HPE BSM 9.25, you must install the HPE BSM 9.25 hotfix. Contact support for more information.

On the HPE OMi machine:

Follow the steps:

- a. Navigate to
On BSM, **Admin > Platform > Users and Permissions > Authentication Management**.
On OMi, **Administration > Users > Authentication Management**.
- b. In the **Single Sign-On Configuration** section, select and copy the value of **Token Creation Key (initString)** field. This **InitString** is used to configure HPE Cloud Optimizer for SSO.

On the HPE Cloud Optimizer machine:

Follow the steps:

- a. Log on as a root user.
- b. Go to `lwssofmconf.xml` file at the following location:
`/opt/OV/www/webapps/PV/WEB-INF/classes/`
- c. Edit the values of `<domain>` and `<initString>` in the `lwssofmconf.xml` file.

For example:

```
<validation>
<in-ui-lwssso>
<lwssoValidation id="ID000001">
<domain><hp.com></domain>
<crypto cipherType="symmetricBlockCipher"
engineName="AES" paddingModeName="CBC" keySize="256"
encodingMode="Base64Url"
initString=<mention the init string here></crypto>
</lwssoValidation>
</in-ui-lwssso>
</validation>
```

Make sure the `<domain>` and `<initString>` are the same as mentioned in the HPE SSO configuration.

After configuring the prerequisites, follow these steps to view the performance graphs in HPE OMi:

1. Navigate to
On BSM, **Applications > Operations Management > Performance Perspective**.
On OMi, **Workspace > Operations Console > Performance Perspective**.
2. In the **View Explorer** pane, select **vPV_Infrastructure**.
3. Select a Configuration Item (CI) to draw the performance graphs.

Customizing the Severity of Alerts

To customize the severity of alerts that are being forwarded to HPOM/OMi, perform the following steps:

On HPOM/OMi machine:

1. Navigate to
On HPOM, **Policy Management > Policy Groups > vPV-Integration**.
On OMi, **<> Administration > Monitoring > Policy Templates > Open Message Interface**.
2. In the right-pane, double click the **vPV-OMIntegration** policy.
The **vPV-OMIntegration** policy edit window appears.
3. Double-click the rule for which you want to customize the Severity.
The rule edit window appears.
4. Set the Severity for the rule as Normal, Minor, Critical, Warning, or Major.
5. Click **OK**.

Integrating HPE Cloud Optimizer with HPE OneView

You can now integrate HPE Cloud Optimizer with HPE OneView. HPE OneView is a converged infrastructure management product that provides a single integrated platform to manage your physical infrastructure. HPE OneView provides simple lifecycle management for enterprise IT—servers, networking, software, power and cooling, and storage.

HPE Cloud Optimizer and HPE OneView integration provides you with the following features:

- A Unified Enclosure View to depict how the physical and virtual infrastructures are connected together and the ability to rapidly identify and resolve potential problems
- View Enclosure and Blade server topology information in Treemap
- Out-of-the-box reports to analyze the risk of shutting down an Enclosure on the

VMware cluster configured on it

- View the CPU, Power, and Temperature utilization metrics in Workbench

HPE Cloud Optimizer and OneView integration allows you to analyze your physical environment from the virtualization perspective. Enclosures monitored by HPE OneView are physical structures that contain server blades, infrastructure hardware, and interconnects. With HPE Cloud Optimizer and HPE OneView integration, the server blades in an enclosure are mapped to the respective hypervisor hosts in HPE Cloud Optimizer. This mapping enables you to visualize the impact of any change in the physical infrastructure on the virtual domain. You can assess the impact of an Enclosure maintenance on the VMware cluster configured on the Enclosure.

For example, a specific c7000 enclosure is expected to be unavailable for a certain duration. With HPE Cloud Optimizer and HPE OneView integration, you can quickly identify the hosts and VMs that would get affected with the enclosure unavailability. The affected hosts and VMs can then be suitably placed to minimize the impact.

Note: You can integrate HPE Cloud Optimizer with HPE OneView 1.20 and 2.00. For the latest support matrix information, see the HPE Software Product Support Matrix.

To integrate HPE Cloud Optimizer with HPE OneView, follow these steps:

On HPE Cloud Optimizer machine:

1. Log on as an administrator.
2. From the HPE Cloud Optimizer home page, click **Settings > Collection & Analysis**.
3. Click **Add Datasource** link.
4. Select **OneView** from the Domain drop-down list.
5. Type the IP address or host name of the HPE OneView appliance in the **OneView IP/Host Name** field.
6. Type the user name of the HPE OneView appliance in the **User Name** field.


Note: The user to be configured in HPE Cloud Optimizer should have PUT and POST access to the following URL:
`https://<IP address of HPE OneView appliance>/rest/login-sessions`

7. Type the password of the HPE OneView appliance in the **Password** field.
8. Click **Test Connection**, if you want to validate the credentials provided. The system validates the credentials and displays a message stating whether the credentials are valid or not. The system also validates if the connection is successful or not.
9. Click **Add**. The Connections list displays the added HPE OneView appliance.

Note: If there is any change in HPE OneView setup, such as addition or deletion of Enclosures and Blade Servers, HPE Cloud Optimizer recommends restarting the data collection for the specific HPE OneView setup. This ensures that the Treemap is updated with the latest changes in HPE OneView setup.

Using HPE Cloud Optimizer features with HPE OneView Integration

After the integration with HPE OneView, HPE Cloud Optimizer features are enhanced with the following capabilities:

Feature	Description
Unified Enclosure View	<p>When HPE Cloud Optimizer is integrated with HPE OneView, the Unified Enclosure View provides a visual representation of the virtual infrastructure overlaid on the actual physical layout of the HPE OneView Enclosure.</p> <p>For more information on the Unified Enclosure View, see the Unified Enclosure View section.</p>
Treemap	<p>When HPE Cloud Optimizer is integrated with HPE OneView, HPE OneView appears as a domain in the Treemap.</p> <p>Select OneView as the domain in the Domain list of Treemap. Treemap appears with the performance information about HPE OneView components.</p>  <p>When OneView is selected as the domain, you can perform the following tasks from the Treemap:</p>

Feature	Description
	<ul style="list-style-type: none"> • View Enclosures and Blade Server topology information • Access Workbench to view the performance metrics, graphs, and reports for Enclosures and Blade Servers. For more information on Enclosure and Blade Server reports, see the Using HPE Cloud Optimizer reports for HPE OneView section. <p>For more information on Treemap, see the Monitoring and Analyzing Environment using Treemap section.</p>
Workbench	<p>To view the performance graphs and reports for Enclosures and Blade Servers monitored by HPE OneView, you can launch the Workbench feature from the Treemap.</p> <p>In the Treemap, right-click an Enclosure or Blade Server and select Launch Workbench. The Workbench page appears with the selected Enclosure or Blade Server displayed in the Configuration Items list.</p> <p>You can perform the following tasks from the Workbench:</p> <ul style="list-style-type: none"> • View the performance metrics for Enclosures and Blade Servers • Draw and view the graphs with the help of available metrics. Additionally, you can save the frequently used graphs as favorites. • View reports to help you analyze the performance of Enclosures and Blade Servers. <p>For more information on Workbench, see the Performance Graphs section.</p>

Unified Enclosure View

The **Unified Enclosure View** is a unique and intuitive representation of the virtual infrastructure overlaid on the actual physical layout of the Enclosures and Server Hardware managed by HPE OneView.

HPE Cloud Optimizer is a unified performance monitoring and capacity analysis tool. HPE Cloud Optimizer has the ability to show resource usage trends and forecast for datacenters, clusters, hosts, and guests. HPE Cloud Optimizer also has the capability to model future resource requirements and analyze the adequateness of existing resources to fulfill future needs. On the other hand, HPE OneView is a converged infrastructure management product that provides a single integrated platform to manage your physical infrastructure.

The Unified Enclosure View allows you to leverage the power of both HPE Cloud Optimizer and HPE OneView to achieve end to end capacity planning and provisioning. HPE Cloud Optimizer and HPE OneView come together to automate day to day management and provisioning tasks. It helps you to intuitively respond to business demands as they occur, resolve potential infrastructure issues before they impact service.

The Unified Enclosure View allows you to perform the following tasks:

- Get a comprehensive virtualized view of the selected Enclosure.
- Visualize the physical and virtual connection between Server Hardware and Hypervisor and Non-Hypervisor hosts.
- Assess the capacity of the Server Hardware by analyzing the Days To Capacity (DTC) value of the host or of the cluster to which the host belongs.
- Analyze the impact of removing a Server Hardware on the capacity of the cluster to which it belongs.
- Perform risk assessment on the sustainability of the cluster on removing an Enclosure.
- View the list of Business Groups having footprints on the hosts configured on the Enclosure.

Navigation













To access the feature, select **Overview**. From the drop-down, select **Unified Enclosure View**.


















The screenshot displays the HPE Cloud Optimizer interface for the 'Gringotts-E1' enclosure. The top navigation bar includes 'Overview', 'Performance', and 'Capacity'. A callout points to the 'Overview' dropdown, stating: 'Launch HP OneView from Physical Layout or Server Hardware view.' The main area shows a grid of 16 host slots. Slots 1, 2, 3, 5, and 11 are populated with host information, including icons for 'Mortgage', 'Personal', and 'Business' groups, and a 'Days To Capacity' value of '>90'. A callout on the left states: 'Get a comprehensive virtualized view of the selected Enclosure.' Another callout on the left states: 'Visualize the Physical and Virtual association between Blade servers and Hosts.' A callout on the right states: 'Analyze the impact of removing a host on the capacity of the cluster to which it belongs.' A callout on the bottom right states: 'Perform risk assessment on the sustainability of the cluster on removing an Enclosure or Server Hardware.' Below the grid, there are buttons for 'Click on a Bay to select/deselect for removal', 'Select/Deselect all', 'Run Scenario', and 'Cancel'. At the bottom, there are tabs for 'Clusters (3)', 'Hosts (5)', and 'Business Groups (4)'. A checkbox 'Show all hosts for clusters' is checked. Below these tabs is a table with the following data:






Cluster Name	Days To Capacity	Actions	Host Name	Blade Server	Enclosure
Mortgage	>90	[Icons]	16.184.42.72	bay 2	Gringotts-E1

Using Unified Enclosure View


The following table provides detailed explanation about using the various features of the Unified Enclosure View:




Task	Description										
Get a comprehensive virtualized front view of the Enclosure.	<p>To get a virtualized view of the Enclosure, select the Enclosure from the OneView tree on the left side. The virtualized front view of the Enclosure appears on the right side.</p> <p>Each box represents an Enclosure Bay. You can view the Bay number on the upper-left corner of the box.</p> <p>The HPE OneView Enclosure comprises of Server Hardware. The Unified Enclosure View gives the following details about the configuration of Server Hardware:</p> <ul style="list-style-type: none"> • Cluster Name, if the specific Server Hardware belongs to a cluster • Name of the Standalone hosts configured on the Server Hardware • Server Hardware that is not being monitored by HPE Cloud Optimizer • Un-configured Server Hardware • Vacant Bays or Bays with no Server Hardware installed <p>The configuration of each Bay is represented uniquely using the following icons:</p> <table> <tr> <th>Icon</th><th>Description</th></tr> <tr> <td></td><td>Server Hardware belongs to a cluster.</td></tr> <tr> <td></td><td>Server Hardware or Host that does not belong to a cluster.</td></tr> <tr> <td></td><td>Non-virtualized Physical Host.</td></tr> <tr> <td></td><td> <p>Server Hardware not being monitored by HPE Cloud Optimizer. The Server Hardware can be a virtualized Host or a non-virtualized Physical Host.</p> <p>To start viewing the virtualization information for the host on the Unified Enclosure View, add the required target to HPE</p> </td></tr> </table>	Icon	Description		Server Hardware belongs to a cluster.		Server Hardware or Host that does not belong to a cluster.		Non-virtualized Physical Host.		<p>Server Hardware not being monitored by HPE Cloud Optimizer. The Server Hardware can be a virtualized Host or a non-virtualized Physical Host.</p> <p>To start viewing the virtualization information for the host on the Unified Enclosure View, add the required target to HPE</p>
Icon	Description										
	Server Hardware belongs to a cluster.										
	Server Hardware or Host that does not belong to a cluster.										
	Non-virtualized Physical Host.										
	<p>Server Hardware not being monitored by HPE Cloud Optimizer. The Server Hardware can be a virtualized Host or a non-virtualized Physical Host.</p> <p>To start viewing the virtualization information for the host on the Unified Enclosure View, add the required target to HPE</p>										





Task	Description								
	<table> <tr> <td></td><td>Cloud Optimizer. For more information on how to add a target, see Adding Datasources.</td></tr> <tr> <td></td><td>Bays with un-configured Server Hardware.</td></tr> <tr> <td>Empty</td><td>Vacant Bays</td></tr> </table>		Cloud Optimizer. For more information on how to add a target, see Adding Datasources .		Bays with un-configured Server Hardware.	Empty	Vacant Bays		
	Cloud Optimizer. For more information on how to add a target, see Adding Datasources .								
	Bays with un-configured Server Hardware.								
Empty	Vacant Bays								
View Days To Capacity (DTC) value of the Server Hardware or of the cluster to which the Server Hardware belongs	<p>The DTC value of the Server Hardware or the cluster to which the Server Hardware belongs to is displayed on the respective Blade Server. The DTC prediction is an estimate of the number of days in which the entity reaches the maximum capacity.</p> <p>If the DTC value of the configured entity is less than 90 days, you can see the following icons that represent the reason for the shortfall.</p> <table> <tr> <th>Icon</th><th>Resource</th></tr> <tr> <td></td><td>CPU</td></tr> <tr> <td></td><td>Memory</td></tr> <tr> <td></td><td>Storage</td></tr> </table> <p>Note: If the Server Hardware belongs to a cluster, the capacity information displayed is for the cluster and not for the respective host.</p>	Icon	Resource		CPU		Memory		Storage
Icon	Resource								
	CPU								
	Memory								
	Storage								
Launch HPE OneView	<p>You can launch HPE OneView from the Unified Enclosure View with an Enclosure view or the Server Hardware view.</p> <ul style="list-style-type: none"> To launch HPE OneView with an Enclosure view, click the  Launch Enclosure in OneView icon besides the Enclosure name. To launch HPE OneView with a Server Hardware view, mouse over on the Server Hardware and click the  Launch Enclosure in OneView icon. 								
View the	To view the CPU, Memory, and Disk utilization of a specific Server								




Task	Description								
resource utilization of the Server Hardware	<p>Hardware, click on the respective Server Hardware. The view flips and displays a bar chart with the CPU, Memory, and Disk utilization values of the Server Hardware.</p> <p>Click the  Flip All icon to view the CPU, Memory, and Disk utilization of all the Server Hardware monitored by HPE Cloud Optimizer.</p>								
Customize the Color Range for DTC values.	<p>The Server Hardwares are highlighted in various colors based on the DTC values. You can customize the color range for different DTC values.</p> <p>To customize the Color Range for DTC values:</p> <ol style="list-style-type: none"> 1. Click the  View/Edit Color Range icon in the upper-right corner of the Unified Enclosure View page. The Color Range for Days to Capacity Values window appears. 2. Set the number of days for Red, Yellow, and Green color codes. The Red, Yellow, and Green color codes represent the High, Medium, and Low risk categories in terms of DTC. 								
View details of all the clusters configured on the Enclosure.	<p>The Cluster tab in the table at the bottom of the page provides details about all the clusters that contain the hosts configured on the Enclosure. The Cluster tab provides the following information:</p> <table border="1"> <thead> <tr> <th>Field</th><th>Description</th></tr> </thead> <tbody> <tr> <td>Cluster Name</td><td>Name of the Cluster.</td></tr> <tr> <td>Days To Capacity</td><td> <p>The number of days the cluster can sustain itself before reaching its maximum capacity.</p> <p>Click the DTC value to open the Cluster Forecast page. For more information on the Forecast, see Forecast.</p> </td></tr> <tr> <td>Actions</td><td> <p>View the Cluster-Enclosure relationship, Capacity Usage Trends, Workbench, and Forecast for the cluster.</p> <p>Click the Cluster-Enclosure relationship icon </p> </td></tr> </tbody> </table>	Field	Description	Cluster Name	Name of the Cluster.	Days To Capacity	<p>The number of days the cluster can sustain itself before reaching its maximum capacity.</p> <p>Click the DTC value to open the Cluster Forecast page. For more information on the Forecast, see Forecast.</p>	Actions	<p>View the Cluster-Enclosure relationship, Capacity Usage Trends, Workbench, and Forecast for the cluster.</p> <p>Click the Cluster-Enclosure relationship icon </p>
Field	Description								
Cluster Name	Name of the Cluster.								
Days To Capacity	<p>The number of days the cluster can sustain itself before reaching its maximum capacity.</p> <p>Click the DTC value to open the Cluster Forecast page. For more information on the Forecast, see Forecast.</p>								
Actions	<p>View the Cluster-Enclosure relationship, Capacity Usage Trends, Workbench, and Forecast for the cluster.</p> <p>Click the Cluster-Enclosure relationship icon </p>								

Task	Description															
		to view the Risk Analysis page. The Risk Analysis page provides a risk analysis of shutting down the Enclosure on the cluster. For more information on the Risk Analysis page, see Risk Analysis .														
	Host Name	Hosts that belong to the specific cluster. By default, all the hosts that are a part of the cluster are displayed. Note: If you clear the Show all hosts for clusters check-box, hosts that do not belong to the selected Enclosure are hidden.														
	Blade Server	The Bay number on which the host is configured.														
	Enclosure	The Enclosure on which the host is configured.														
View details of all the Hosts configured on the Enclosure.	<p>The Hosts tab in the table at the bottom of the page provides details of all the hosts or Server Hardware configured on the Enclosure. The Hosts tab provides the following information:</p> <table><tr><th>Field</th><th>Description</th></tr><tr><td>Host Name</td><td>Host Name or IP address of the Server Hardware.</td></tr><tr><td>Days To Capacity</td><td>The number of days the host can sustain itself before reaching its maximum capacity. Click the DTC value to open the Host Forecast page. For more information on the Forecast, see Forecast.</td></tr><tr><td>Cluster Name</td><td>Name of the cluster to which the host belongs.</td></tr><tr><td>Blade Server</td><td>The Bay number on which the Host is configured in the Enclosure.</td></tr><tr><td>Virtualization Domain</td><td>Virtualization domain of the Host.</td></tr><tr><td>Actions</td><td>View the Capacity Usage Trends, Workbench, and Forecast for the Host.</td></tr></table>		Field	Description	Host Name	Host Name or IP address of the Server Hardware.	Days To Capacity	The number of days the host can sustain itself before reaching its maximum capacity. Click the DTC value to open the Host Forecast page. For more information on the Forecast, see Forecast .	Cluster Name	Name of the cluster to which the host belongs.	Blade Server	The Bay number on which the Host is configured in the Enclosure.	Virtualization Domain	Virtualization domain of the Host.	Actions	View the Capacity Usage Trends, Workbench, and Forecast for the Host.
Field	Description															
Host Name	Host Name or IP address of the Server Hardware.															
Days To Capacity	The number of days the host can sustain itself before reaching its maximum capacity. Click the DTC value to open the Host Forecast page. For more information on the Forecast, see Forecast .															
Cluster Name	Name of the cluster to which the host belongs.															
Blade Server	The Bay number on which the Host is configured in the Enclosure.															
Virtualization Domain	Virtualization domain of the Host.															
Actions	View the Capacity Usage Trends, Workbench, and Forecast for the Host.															

Task	Description														
View details of the Business Groups having footprints on the Enclosure.	<p>The Business Groups tab in the table at the bottom of the page provides details of all the Business Groups having footprint on the Enclosure.</p> <p>A Business Group appears in this list if it has a presence on the Hosts configured on the Blade Servers or on a Cluster that has a footprint on the Enclosure. Following is a visual representation that shows when does a Business Group appear in the Unified Enclosure View:</p>  <pre> graph LR BG[Business Group A] --> VM1[VM1] BG --> VM2[VM2] BG --> VMn[VM-n] VM1 --> C1[Cluster 1] C1 --> H1[Host 1] C1 --> H2[Host 2] C1 --> H3[Host 3] VM2 --> H4[Host 4] H1 --> E1[Enclosure 1] H2 --> E1 H3 --> E1 H4 --> E1 </pre> <p>The Business Groups tab provides the following information:</p> <table> <thead> <tr> <th>Field</th><th>Description</th></tr> </thead> <tbody> <tr> <td>Business Group</td><td>The Business Group instances having footprints on the Enclosure.</td></tr> <tr> <td>Business Grouping</td><td>Name of the Business Grouping having footprints on the Enclosure.</td></tr> <tr> <td>Type of Members</td><td>The type of members constituting the Business Group. The members can be VMs, Hosts, or Datastores.</td></tr> <tr> <td>Days To Capacity</td><td>The number of days the Business Grouping can sustain itself before reaching its maximum capacity. Click the DTC value to open the Group Forecast page. For more information on the Forecast, see Forecast.</td></tr> <tr> <td>Cluster Name</td><td>Name of the Cluster having footprints of the Business Grouping.</td></tr> <tr> <td>Host Name</td><td>Name of the Host having footprints of the Business Grouping.</td></tr> </tbody> </table>	Field	Description	Business Group	The Business Group instances having footprints on the Enclosure.	Business Grouping	Name of the Business Grouping having footprints on the Enclosure.	Type of Members	The type of members constituting the Business Group. The members can be VMs, Hosts, or Datastores.	Days To Capacity	The number of days the Business Grouping can sustain itself before reaching its maximum capacity. Click the DTC value to open the Group Forecast page. For more information on the Forecast, see Forecast .	Cluster Name	Name of the Cluster having footprints of the Business Grouping.	Host Name	Name of the Host having footprints of the Business Grouping.
Field	Description														
Business Group	The Business Group instances having footprints on the Enclosure.														
Business Grouping	Name of the Business Grouping having footprints on the Enclosure.														
Type of Members	The type of members constituting the Business Group. The members can be VMs, Hosts, or Datastores.														
Days To Capacity	The number of days the Business Grouping can sustain itself before reaching its maximum capacity. Click the DTC value to open the Group Forecast page. For more information on the Forecast, see Forecast .														
Cluster Name	Name of the Cluster having footprints of the Business Grouping.														
Host Name	Name of the Host having footprints of the Business Grouping.														

Task	Description		
	<table border="1"> <tr> <td data-bbox="414 283 771 367">Blade Server</td><td data-bbox="771 283 1404 367">The Bay number on which the host is configured.</td></tr> </table>	Blade Server	The Bay number on which the host is configured.
Blade Server	The Bay number on which the host is configured.		
Impact Analysis of removing a Server Hardware on the Cluster.	<p>You can analyze the impact of removing a Server Hardware on the Cluster.</p> <p>To model removing a Server Hardware from the enclosure, perform the following steps:</p> <ol style="list-style-type: none"> 1. Click the Model Enclosure Maintenance button. All the Server Hardware that belong to a cluster are highlighted with a selection  icon. 2. Click the selection  icon to select a Bay for removal. Click the selection icon again if you want to deselect. 3. Click the calendar  icon on the upper-right corner of the Bay representation. The Select Dates pop-up window appears. 4. Select a Start Date and End Date for which you want to model removing the Server Hardware from the Enclosure. 5. Click Run Scenario. The Cluster table at the bottom of the page is updated to reflect the impact of removing a host on the respective cluster. The DTC value of the cluster is adjusted to accommodate the removal of the host. 6. Click the DTC value to open the Capacity Modeler page. In the Capacity Modeler page, you can view the impact of removing a host on the CPU, Memory, and Storage resources of the cluster. If the DTC value for CPU or Memory is less than 90 days, you can find additional capacity in un-configured Server Hardware using the Find Capacity feature. For more information about Capacity Modeler and Find Capacity, see Capacity Modeler and Find Capacity respectively. 		
Risk Analysis of shutting down an Enclosure on the Cluster.	<p>You can assess if a cluster can sustain itself when one of the enclosures on which the cluster is configured is shut down for maintenance.</p> <p>The cluster is considered to be at High, Medium, or Low risk based on the DTC values configured in the Color Range for DTC Values feature. For more details on setting the DTC values for High, Medium, and Low risk, see Customize the Color Range for DTC values.</p>		

Icon	Description
	Server Hardware belongs to a cluster.
	Server Hardware or Host that does not belong to a cluster.
	Non-virtualized Physical Host.
	Bays with un-configured Server Hardware.
Empty	Vacant Bays

Icon	Resource
	CPU
	Memory
	Storage

Field	Description
Host Name	Host Name or IP address of the Server Hardware.
Days To Capacity	The number of days the host can sustain itself before reaching its maximum capacity. Click the DTC value to open the Host Forecast page. For more information on the Forecast, see Forecast .
Cluster Name	Name of the cluster to which the host belongs.
Blade Server	The Bay number on which the Host is configured in the Enclosure.
Virtualization Domain	Virtualization domain of the Host.

Actions	View the Capacity Usage Trends, Workbench, and Forecast for the Host.
---------	---

Field	Description
Business Group	The Business Group instances having footprints on the Enclosure.
Business Grouping	Name of the Business Grouping having footprints on the Enclosure.
Days To Capacity	The number of days the Business Grouping can sustain itself before reaching its maximum capacity. Click the DTC value to open the Business Group Forecast page. For more information on the Forecast, see Forecast .
Cluster Name	Name of the Cluster having footprints of the Business Grouping.
Host Name	Name of the Host having footprints of the Business Grouping.
Blade Server	The Bay number on which the host is configured.

Working with Business Groups

Business Grouping in Cloud Optimizer allows you to customize the operational view of your virtual infrastructure, based on your business applications and services. Using Business Grouping, you can categorize and group your Servers, Hypervisors, or Datastores based on the needs and priorities of your organization. Servers include Physical Servers and VMs. Hypervisors are Hypervisor host machines of VMWARE, Hyper-V, KVM, or Xen domains. A Business Group is a logical grouping of entities representing a common business requirement.

Note: You can create business groups for VMware, KVM, Xen, and Hyper-V domains.

Business Groups are closely integrated with all the features of HPE Cloud Optimizer, enabling you to perform the following:

- Ability to work with a restricted set of virtual elements in a well-defined group
- Define and create your own set of metrics for your Business Group

- View user defined or custom business metrics such as response time, number of transactions and so on, for the entities in a Business Group
- Analyze the capacity, usage, and allocation trends for various resources in a Business Group
- Forecast the resource utilization trends of entities in a Business Group

Note: While calculating the total capacity of a Business Group of Servers, the capacity of powered off and suspended members in the group is also considered.

Following are some use-case scenarios to understand the usability of Business Grouping in the real-time virtualized environments.

- John is a virtualization Subject Matter Expert (SME) for an organization having three departments - Finance, Marketing, and Human Resource. John wants to group all the Servers, Hypervisors, and Datastores in his virtualized infrastructure based on the departments they have been allocated to. Using HPE Cloud Optimizer, John can achieve this by creating a Business Grouping for these groups. This would enable him to monitor and analyze the performance and resource allocation of the individual departments efficiently.
- Bill is a virtualization Subject Matter Expert (SME) in a telecommunication service provider firm. His responsibilities include continuous monitoring of the performance of digital services being offered in a specific region. Bill has created a Business Grouping of all the Servers that provide mobile SIM activation service in his region. Bill's organization has announced a promotional offer for a week and he is expecting the SIM activation requests to double during that time period. During this period, Bill has to ensure zero downtime and excellent connectivity. He needs to identify and resolve potential performance bottlenecks in the infrastructure. Bill can use the forecasting capability of Business Grouping and ensure that the virtual infrastructure is capable of handling the increase in SIM activation requests.

Creating a Business Grouping

HPE Cloud Optimizer provides an easy to use and flexible wizard to create Business Groupings. The wizard guides you through the steps needed to create a Business Grouping. The wizard helps you configure the basic settings required for creating a new Business Grouping.

You can define metrics, create multiple group instances, add members, define Treemap views, and view summary of the Business Grouping.

Navigation

To access Business Group from the HPE Cloud Optimizer home page, select **Settings >**


Business Group.

Using the Business Grouping Creation Wizard



To create a Business Grouping, follow these steps:

Step	Task	Description
1	Assigning Name and Type	In the Name and Type tab, enter Business Grouping name and select a type for the Business Grouping. You can also add a small description of the Business Grouping. A Business Grouping can contain Servers, Hypervisors, or Datastores. Click Next .
2	Configuring a Group	<p>In the Group Configuration tab, you can upload the CSV file, which contains Group names, User-defined metric names, and the Values for the user-defined metrics. You can also enter the Group names manually.</p> <p>Note: The User-defined metrics can be uploaded only through CSV. Use the Upload CSV option to create User-defined metrics.</p> <p>To upload the CSV, perform the following steps:</p> <ol style="list-style-type: none"> 1. Select the Upload CSV option. 2. Click Choose File to browse and select the CSV file. 3. Click Next. <p>HPE Cloud Optimizer reads the CSV file at each data collection interval irrespective of the file being updated or not.</p> <p>Note: Click Download Sample CSV to download a sample CSV file. You can study the sample CSV file to learn how to create your own CSV files or you can add your own data to the sample CSV file and use it in the wizard. The data in the sample CSV files is provided to help you understand the type of values that you can enter in each column. The sample data is not intended to be</p>

Step	Task	Description
		<p>used in the wizard.</p> <p>To enter the Group names manually, perform the following steps:</p> <ol style="list-style-type: none"> 1. In the Enter Group Name text box, enter the group name. 2. Click Add. 3. Repeat steps 1 and 2 to add multiple groups.
3	Adding Members	<p>In the Members tab, you can add members using a filter or by uploading a CSV file, which contains Group names and Member names.</p> <p>To add members for the Group using filter condition, perform the following steps:</p> <ol style="list-style-type: none"> 1. For Type of Members as Servers, in the Member Filter section, select Name, Annotation, IP Address, or Host Name from the Servers with drop-down and enter a value (matching pattern) in the matching box that you would like to search in all Servers. <p>While filtering members based on Annotation, HPE Cloud Optimizer searches in the Annotations > Notes section for the entered pattern.</p>  <p>For Type of Members as Datastores or Hypervisors, enter the matching name in the Enter Pattern box.</p> <ol style="list-style-type: none"> 2. Click Next. A wildcard search is performed on all the Server, Hypervisor, or Datastore names. All the entities matching the search criteria are added to the Group instance. <p>To add members for the Groups using a CSV file, perform the following steps:</p> <ol style="list-style-type: none"> 1. Click Browse.


Step	Task	Description
		<ol style="list-style-type: none"> 2. Select the CSV file. 3. Click Next. <p>Note: Click Download Sample CSV to download a sample CSV file. You can study the sample CSV file to learn how to create your own CSV files or you can add your own data to the sample CSV file and use them in the wizard. The data in the sample CSV file is provided to help you understand the type of values that you can enter in each column. The sample data is not intended to be used in the wizard.</p>
4	Selecting or Adding Predefined Metrics	<p>In the Pre-defined Metrics tab, you can select and add metrics from the list of predefined HPE Cloud Optimizer metrics. However, some predefined metrics are added by default on the basis of the type selected for the Business Grouping.</p> <p>To add a predefined metric, perform the following steps:</p> <ol style="list-style-type: none"> 1. In the Metric column, select a metric from the drop-down list that you would like to add for the Business Grouping. 2. Click Add to add multiple metrics. 3. Click Next.
5	Selecting User-defined Metrics	<p>In the User Defined Metrics tab, you can select the user-defined metrics that you would like to enable for the Business Grouping. The table lists all the metric names defined in the CSV file. You can also select a maximum of two metrics to be available for forecasting.</p>
6	Defining Treemap View	<p>In the Treemap Configuration tab, you can define two Treemap views. The second Treemap view is optional.</p> <p>To create Treemap views, perform the following steps:</p> <ol style="list-style-type: none"> 1. In the Name field, type a name for the Treemap view. 2. In the Min and Max fields, type minimum and maximum values. These values are used to filter the information displayed in the Treemap. 3. In the table, select Size Metric and Color Metric for the Treemap view.


Step	Task	Description
7	Viewing Summary	In the Summary tab, view a summary of the Business Group. You can go back to any of the earlier tabs and change the configuration.
8	Saving and Deploying the Business Group	Click Deploy to save and create the Business Grouping.

Using Business Groups with other HPE Cloud Optimizer Features

The Business Grouping feature appears with several other HPE Cloud Optimizer features. The following table explains the functioning of Business Groups with other HPE Cloud Optimizer features:

Feature	Description
Treemap	<p>After a Business Grouping is deployed, it appears in the Treemap as a domain. You can select the Business Grouping name as the domain and do an analysis of the performance of the Groups in a Business Grouping.</p> <p>For example, if you create a Business Grouping as Department, it appears in the Domain list of the Treemap.</p>

Feature	Description
	 <p>Note: HPE Cloud Optimizer recommends refreshing the Treemap after a new Business Grouping is deployed. A newly deployed Business Grouping may take around five minutes to appear in the Treemap.</p> <p>When a Business Grouping is selected as the domain, you can perform the following tasks from the Treemap:</p> <ul style="list-style-type: none"> • View the number of Groups and type of members in the Resource pane. • View the Groups based on the Treemap View created while defining the Business Grouping definition. • Drill down to view the performance of Servers, Hypervisors, or Datastores in a Group. • Right-click and select Forecast to view forecast for predefined metrics and user-defined metrics of the Group. <p>For more information on Treemap, see the Monitoring and Analyzing Environment using Treemap section.</p>
Integration with HPE Cloud Service Automation (HPE CSA)	<p>When HPE Cloud Optimizer is integrated with HPE CSA, you can create a CSA Business Grouping. When CSA Business Grouping is selected as the domain in Treemap, you can view HPE CSA organizations as Groups and HPE CSA subscriptions as Servers.</p> <p>Creating CSA Business Grouping</p> <p>To create CSA Business Grouping, perform the following steps:</p> <ol style="list-style-type: none"> 1. Log on to HPE Cloud Optimizer as an administrator. 2. From the HPE Cloud Optimizer home page, select Settings

Feature	Description
	<p>> Integration.</p> <ol style="list-style-type: none"> 3. In the HPE Cloud Service Automation (CSA) section, click Edit. The CSA Integration Settings dialog box appears. 4. In the Edit CSA Credentials dialog box, perform the following tasks: <ol style="list-style-type: none"> a. Type the URL for which you want to integrate CSA with HPE Cloud Optimizer in the CSA URL field. For example: <code>https://<CSA Hostname or IP address>:8444</code> b. Type the administrator user name and password in the User Name and Password field respectively. c. The Deploy Business Group option is selected by default. If the Deploy Business Group option is selected, a Business Grouping called CSA is deployed in HPE Cloud Optimizer.  <ol style="list-style-type: none"> d. In the Server Component Types box, type the name of the custom server component type. Select the Exact Match check box if you know the exact value for Server Component Types. e. In the Host Name Property box, type the name of the custom property containing the VM name.

Feature	Description
	<p>Note: CSA has the option of predefined and custom workflow. For a predefined workflow the values for Server Component Types and Host Name Property are optional. For a custom workflow you must enter the values for Server Component Types¹ and Host Name Property².</p> <p>f. Click Save. A message stating that the CSA information is saved successfully appears.</p> <p>Note: If you change any configuration property and would like to save the CSA information again, it is highly recommended to reset the CSA users information using the Reset CSA Users Info button before saving.</p> <p>After the configuration, CSA is integrated with HPE Cloud Optimizer. A CSA Business Grouping is deployed in HPE Cloud Optimizer and it appears in the Domain list of the Treemap.</p> <p>Removing CSA Business Grouping</p> <p>To remove CSA Business Grouping, perform the following steps:</p> <ol style="list-style-type: none"> 1. Log on as administrator. 2. From the HPE Cloud Optimizer home page, select Settings > Business Group. 3. In the List of Groupings table, click the Delete icon for the CSA Business Grouping. <p>For more information on HPE Cloud Optimizer and HPE CSA integration, see the Integrating HPE Cloud Optimizer with HPE CSA section.</p> <p>For more information on Treemap, see the Monitoring and Analyzing Environment using Treemap section.</p>
Workbench	To view the performance graphs for the Business Grouping, you can launch the Workbench feature from the Groups or the Group

¹Name of the custom server component type.²Custom property containing the VM name.

Feature	Description
	<p>members in the Treemap.</p> <p>In the Treemap, right-click a Group or a Group member and select Launch Workbench. The Workbench page appears with the selected Group or a Group member name in the Configuration Items list.</p> <p>You can perform the following tasks from the Workbench:</p> <ul style="list-style-type: none"> • View the user-defined and predefined metrics. • Visualize and triage the performance of Groups and Group members. • Draw and view graphs with the help of available metrics. Additionally, you can save frequently used graphs as favorites. <p>For more information on Workbench, see the Performance Graphs section.</p>
Forecast	<p>To view forecast for Groups and Group members, you can launch the Forecast feature from Treemap.</p> <p>In the Treemap, right-click a Group or a Group member and select Forecast. The Forecast page appears for the selected Group or a Group member.</p> <p>You can perform the following tasks from the Forecast page:</p> <ul style="list-style-type: none"> • View the forecast for predefined metrics (in percentage). • View the forecast for user-defined metrics that were selected while creating the Business Grouping definition (in units). • View the information on Day To Capacity (DTC) and total capacity for the respective Group members. <p>For more information on Forecast, see the Forecast section.</p>

Resource Meter

The Resource Meter adds another dimension to HPE Cloud Optimizer's strong capacity planning and management capabilities. You can use the Resource Meter to track the CPU, Memory, and Storage allocation of a set of VMs over a specific time interval. It also enables us to charge back the respective groups, which are using the services. The Resource Meter in HPE Cloud Optimizer allows you to calculate the billing cost for the CPU, Memory, and Storage resources used by your business applications and services.

It gives a quick view of how the virtual machines are being used over time and helps you gain visibility into the true costs of the virtualized workloads. It also gives the potential savings or expense on a given sizing recommendation.


Note: You can view the Resource Meter only if you have created Business Groups for VM's.


Navigation

To access the feature, from the HPE Cloud Optimizer home page, select **Groups > Resource Meter**.


Using the Resource Meter

1. From the left pane, select the Business Group.
2. Select the time period for which you want to calculate the usage. To select the

period, from the Date Range section, click the  icon and select **Today**, **Yesterday**, **Last 7 Days**, **Last 30 Days**, **This Month**, or **Custom Range** as required, and click Apply.



By default, the **Last 7 Days** option is selected. When you select any of these options, it displays the allocations for all the VMs for a Business Group for that period ending in the current date. If you select **Custom Range**, specify the start date

and end date, from the , and click **Apply**. The page refreshes to show the current and recommended cost.

Note:

CPU-Days - It is defined as the product of number of CPUs and number of days. For example, If 4 CPUs are used for first 15 days and 8 CPUs are used for next 15 days in a month, the CPU-Days is calculated as follows: $(4 \times 15) + (8 \times 15) = 180$

Num CPU-Days.

Memory (GB)-Days - It is defined as the product of the amount of memory allocated and number of days. For example, if a VM is running with 4 GB of memory for 5 days, the Memory (GB)-Days is calculated as follows: $5 \times 4 = 20$ Memory (GB)-Days.

Disk (GB)-Days - It is defined as the product of the amount of disk provisioned and number of days. For example, if a VM is running with 20 GB of hard disk for 5 days, the Disk (GB)-Days is calculated as $20 \times 5 = 100$ Disk (GB)-Days.

- Click  **Advanced Settings** on the right-top corner. The Settings page appears.

Note: When you open Resource Meter for the first time, the Settings page opens by default.

- Enter cost per day for CPU, Memory, and Disk in the **CPU**, **Memory (GB)**, and **Disk (GB)** fields.
- Enter the type of currency in the **Type of Currency** field.

Note: In case of data gap in the selected time interval, the cost per day per unit for that interval is calculated based on the available options. The options are as follows:


Zero - When you select **Zero**, the data gap period is not charged.

Before - When you select **Before**, the data gap period is charged as per the value before that interval.

- Click **Save**. The Summary for group page appears with the current cost and recommendation. On mouse over on the Total Cost column, you can view the formula used for the calculation.
- Click **Show Details** to view the consumption of resources and recommendations for the individual VMs. The table contains the following details:

Name	Details
VMName	Name of the Virtual Machine
CPU-Days	Number of CPUs number of days
Memory (GB)-Days	Total memory used

Name	Details
Disk (GB)-Days	Total hard disk used
CurrentCPU	Current CPU usage
RecommendedCPU	Recommendation for CPU usage
CurrentMemory	Current memory usage
Recommended memory	Recommendation for memory usage
Actions	View Capacity Usage Trends , Workbench , and VM Forecast for a VM.

Note: The dimmed  VM icon in the Details table represents powered off VMs. Also, you can ignore if multiple VMs of same name are available in the Details page. These VMs were created at some point of time in the past and deleted. On mouse over, you can view the start date and end date for these VMs.

8. Click  to take a print out of the Summary for group page.

Business Metric Analyzer

The **Business Metric Analyzer** adds another dimension to HPE Cloud Optimizer's strong business grouping functionality. The Business Metric Analyzer allows you to use Business Groups to effectively plan and manage capacity allocation based on changing business requirements.

With the Business Metric Analyzer, you can correlate the user-defined business metrics of a group with the pre-defined infrastructure performance and utilization metrics. With this correlation, you can analyze the impact of change in business metric values on the performance and utilization metrics. Based on this predictive analysis, the Business Metric Analyzer also provides appropriate CPU and Memory recommendations for the virtualized entities in the Business Group.

Note: While calculating the total capacity of a Business Group of VMs, the capacity of powered off and suspended members in the group is also considered. Therefore, the CPU and Memory short-fall predictions by the Business Metric Analyzer are based on the total capacity of the group including the powered off and the suspended VMs.

Navigation

To access the feature, select **Groups**. From the drop-down, select **Business Metric Analyzer**.

Note: The Groups tab is enabled only if you have created Business Groups.

Using the Business Metric Analyzer

1. From the left pane, select the Business Group for which you would like to do the metric analysis.

If the business group contains business metrics, you can see the available business metrics in the right-pane. You can also upload business metrics for the group by uploading a CSV file.

To upload business metrics data using a CSV file, perform the following steps:

- a. Click **Download Sample CSV** to download a sample CSV file. You can study the sample CSV file to learn how to create your own CSV files or you can add your own data to the sample CSV file and use it here. The data in the sample CSV files is provided to help you understand the type of values that you can enter in each column.

Following are the details of the sample CSV:

Field	Description
Timestamp	The Epoch timestamps of the data provided for the business metrics. The Epoch timestamps must be specified in seconds.
Metric1	Name of the first business metric. For example, num_transactions. The column contains the metric values for the first business metric.
Metric2	Name of the second business metric. For example, total_revenue. The column contains the metric values for the second business metric.

Note:

- If you specify more than two metric columns in the CSV, the remaining columns are ignored.
- The data granularity in the CSV file should not be more than the data collection interval of HPE Cloud Optimizer (5 minutes or 15 minutes).

- b. Click **Choose File** to browse and select the CSV file.

The **Available/Uploaded Metrics** list is updated with the metrics available in the uploaded CSV file.

2. Click the metric name for which you want to perform the analysis.

The left box refreshes to display the following information:

- A graph displaying the metric values of the selected metric over a period of time.
- A graph displaying the CPU and Memory utilization for the corresponding metric values.

The X-axis displays the selected metric values. The Y-axis displays the corresponding CPU and Memory utilization values.

3. In the right box, enter the value the selected business metric is expected to reach. Click **Analyze**.

The Business Metric Analyzer analyzes and correlates the value of the business metric with the memory and CPU utilization metrics. You can see the following information:

- Graph in the left box refreshes to show the CPU and memory utilization trends for the Business Group, assuming that the business metric reaches the specified value.

The X-axis displays the minimum and maximum metric values. The maximum value is the value that the business metric is expected to reach. The Y-axis displays the CPU and Memory utilization variation corresponding to the metric values.

- Scale In/Out recommendations for the business metric

You can see the amount of CPU and Memory required by the environment if the business metric reaches the entered value. You can also see the number of reference VMs that should be created to offset the shortfall. The default configuration of reference VM in HPE Cloud Optimizer is 2 vCPU, 4GB Memory, and 40GB storage.

- Scale Up/Down recommendations for VMs

CPU and memory recommendations for the individual VMs of the group.

Use Case

Bill is a virtualization Subject Matter Expert (SME) in a telecommunication service provider firm. His responsibilities include continuous monitoring of the performance of digital services being offered in a specific region. Bill's organization has announced a promotional offer for a week and he is expecting the SIM activation requests to double during that time period. Bill needs to ensure that the virtual infrastructure is capable of

handling the increase in SIM activation requests. This use case demonstrates how Bill can use the Business Metric Analyzer to accomplish his objectives.

The tasks that Bill must perform as an Administrator are as follows:

1. Log on to the HPE Cloud Optimizer.
2. Create a Business Group of all the VMs that provide mobile SIM activation services in his region.
3. Create user-defined business metrics for the Group or upload business metrics using the CSV. For example,
 - Metric to track the number of activations (num_transactions)
 - Metric to track the revenue generated (revenue) with the increasing number of activations

For more information on creating Business Groups, see [Creating a Business Grouping](#).

4. From the HPE Cloud Optimizer home page, select **Groups > Business Metric Analyzer**.
5. In the left pane, select the Business Group of VMs that provide the mobile SIM activation services.

The **Available/Uploaded Metrics** list is updated with the metrics available in the Business Group - num_transactions and revenue.

6. Select the metric **num_transactions**.

The left box refreshes to display the following information:

- A graph displaying the metric values of the selected metric over a period of time.
 - A graph displaying the CPU and Memory utilization for the corresponding metric values.
7. In the right box, enter an expected value for the selected business metric. For example, 2000. This is the estimated number of SIM activation requests expected to come during the period of promotional offer.
 8. Click **Analyze**.

The Business Metric Analyzer analyzes and correlates the value of the metric **num_transactions** with the memory and CPU utilization metrics. You can see the following information:

- Graph that shows the CPU and memory utilization trends for the Business Group, assuming that the business metric reaches the specified value.
- Scale In/Out recommendations for the business metric

You can see the amount of CPU and Memory required by the environment if the business metric reaches the entered value. You can also see the number of reference VMs that should be created to offset the shortfall. The default configuration of reference VM in HPE Cloud Optimizer is 2 vCPU, 4GB Memory, and 40GB storage.

- Scale Up/Down recommendations for VMs
CPU and memory recommendations for the individual VMs of the group.

Based on this analysis, Bill can get valuable insight into the expected behavior of the infrastructure during the promotional offer.

Configuring HPE Cloud Optimizer

The following list specifies the tasks you can perform as an Administrator to get optimum results from HPE Cloud Optimizer:

Note: The following tasks are not applicable for the Archive Extractor version of the product.

- ["Configuring Session Timeout Interval" on page 85](#)
- ["Configuring Secure Client Connect Timeout Interval" on page 86](#)
- ["Configuring Java Virtual Machine Memory" on page 86](#)
- ["Configuring HPE Cloud Optimizer to use Public Key Infrastructure Authentication" on page 87](#)

Conventions Used

The following conventions are used in the sections when referring to the location of files on the HPE Cloud Optimizer server system:

Directory	Description
<install_dir>	The directory where HPE Cloud Optimizer is installed. The default location is <i>/opt/OV</i> . For Archive Extractor version of the product, the directory is where the file is extracted.
<data_dir>	The common data directory where data files and logs related to HPE Software products are stored. The default location is <i>/var/opt/OV</i> . For Archive Extractor version of the product, the location is <i>/data</i> .
<bin_dir>	The directory where the binaries reside. The default location is

Directory	Description
	<code>/opt/OV/bin</code> . For Archive Extractor version of the product, the location is same as the <code><install_dir></code> location.
<code><systemname></code>	<p>The name of the system on which the HPE Cloud Optimizer server is running.</p> <p>When accessing HPE Cloud Optimizer through a web browser, you must replace this variable with the actual system name or IP address of the system.</p>

PVconfig.ini Parameters

The following table lists the parameters, which must be specified in the `PVconfig.ini` file to configure HPE Cloud Optimizer:

Parameter	Description
TRACELEVEL	<p>You can use this parameter to turn on or turn off tracing. When you set this parameter to 1, tracing is turned on and the trace log files are generated.</p> <p>Setting this to 0 turns off tracing. It is recommended to turn on or turn off tracing by using the trace on or trace off commands.</p>
GRAPH_AUTOREFRESH_RATE	You can use this parameter to specify the interval at which HPE Cloud Optimizer refreshes the graphs automatically. The value you specify for this parameter is in seconds. For example, if you specify the value as 120, the graph will be refreshed every two minutes.
DAILY_MAINTENANCE_TIME	You can use this parameter to schedule the maintenance task at a time. The value you specify here is in HH:MM format. For more information, see "Daily Maintenance" on page 94 .
DIAGVIEW_TABLE_ROWCOUNT	You can use this parameter to configure the number of rows in a drill-down table in the Workbench page. The default value is 10000.
SHOW_MOUSE_HOVER_DEFAULT	When you move the mouse on the graph area of a drawn graph from the HPE Cloud Optimizer user interface, a pop-up displaying the actual value of the data point and the time interval of the selected data appears. You can configure this parameter to enable or disable the mouse-hover option. The default value is

Parameter	Description
	<p>TRUE/YES and the pop-up displaying the details of the data appears when you move the mouse on the selected data on the graph. If you set this value to FALSE/NO, the pop-up displaying the value and time of the data point does not appear when you move the mouse over the graph.</p> <p>Note: The value selected from the user interface menu option overrides the setting specified for this parameter.</p>
SECURE_CLIENT_CONNECT_TIMEOUT	<p>You can use this parameter to specify the timeout period for which secure communication HTTPS protocol is used for connecting to the HPE Cloud Optimizer server. The default timeout period is one second. For more information, see "Configuring Secure Client Connect Timeout Interval" on the next page.</p>
RTV_GRAPH_DATAPOINTS	<p>You can use this parameter to configure the maximum number of data points in the graphs drawn using Workbench.</p>
TRUSTED_SITE_URL	<p>You can use this parameter to configure the trusted website for HPE Cloud Optimizer. An application will be able to embed HPE Cloud Optimizer page within itself only if the <i>TRUSTED_SITE_URL</i> is set in the <i>PVconfig.ini</i> file.</p> <p>For example, if the application that wants to embed HPE Cloud Optimizer has the URL <i>mydashboard.company.com/home/</i>, you must do the following configuration in the <i>PVconfig.ini</i> file:</p> <p><i>TRUSTED_SITE_URL=https://<application_URL, for example, mydashboard.company.com></i></p> <p>Note: This restriction is enforced only by the Internet Explorer and Firefox browsers at the time of this release.</p>

Configuring Session Timeout Interval

You can set the expiry time or timeout interval for a user session. If a user session is inactive for the interval you specified, the session expires and the system logs out the user. The default timeout interval for all user sessions is set to one hour.

To modify the default timeout interval, perform the following steps:

1. Open the `PVconfig.ini` file available in the `<data_dir>/conf/perfdirectory`.
2. Specify an interval for the `SESSION_TIMEOUT` parameter in minutes. For example, set it to 30 minutes.
3. Save the file.
4. Restart HPE Cloud Optimizer.

Configuring Secure Client Connect Timeout Interval

As a HPE Cloud Optimizer administrator, you can set the expiration time or timeout interval at which HPE Cloud Optimizer can connect to the server through HTTPS communication protocol. When HPE Cloud Optimizer users use secure communication HTTPS protocol to connect to the server, the specified time-out interval on the secure channel is initially tried. The default timeout interval is set to one second.

To modify the default timeout interval, perform the following steps:

1. Open the `PVconfig.ini` file available in `<data_dir>/conf/perf` directory.
2. Specify an interval for the `SECURE_CLIENT_CONNECT_TIMEOUT` parameter in seconds, depending on the speed of your network connection.

For example, set it to 10 seconds.

3. Save the file.
4. Restart HPE Cloud Optimizer.

Note: When HPE Cloud Optimizer users use regular communication protocol (HTTP) to connect to the server, the secure channel through HTTPS is initially tried for the specified time-out interval and then HTTP is used.

Configuring Java Virtual Machine Memory

To avoid out of memory exceptions caused due to insufficient memory in the Java Virtual Machine (JVM), you can configure HPE Cloud Optimizer to stop accepting requests by using the `JVM_MIN_MEMORY` parameter. By setting this parameter to a specific value, HPE Cloud Optimizer does not accept requests, if the available memory is less than the value specified and displays the error message:

This request cannot be processed, because the HPE Cloud Optimizer server has run out of memory.

If the request is for XML graphs, HPE Cloud Optimizer displays an XML file with the **Out of Memory Error** message. As this parameter is not set by default, you must specify the parameter in the `PVconfig.ini` file.

To specify `JVM_MIN_MEMORY`, perform the following steps:

1. Open the `PVconfig.ini` file available in the `<data_dir>/conf/perf` directory.
2. Enter the `JVM_MIN_MEMORY` parameter and specify a value, for example 5.

The value you enter here are in MBs (megabytes). For example, if you define the memory value as 5 MB and the Java heap size is less than 5 MB, HPE Cloud Optimizer does not accept requests until the available memory for the Java heap size is greater than the value specified for the `JVM_MIN_MEMORY` parameter.

3. Save the file.
4. Restart HPE Cloud Optimizer.

The default value for maximum JVM heap size is 512 MB. If you anticipate a higher load, set the heap size to a higher value.

Configuring HPE Cloud Optimizer to use Public Key Infrastructure Authentication

The Public Key Infrastructure (PKI) is used for the implementation of Common Access Cards (CAC) given to employees of an institution. PKI is used to create, manage, and revoke certificates. CAC enhances the security and access to hardware and software with the inclusion of the digital signature and data encryption technologies used. HPE Cloud Optimizer supports the use of CAC and enables you to map client certificates to user accounts (administrator, customers, and users) and use the Public Key Infrastructure (PKI) for user authentication.

After configuring HPE Cloud Optimizer to use the client certificates, users can log on to HPE Cloud Optimizer with the X.509 certificates without having to manually type in user names and passwords.

Note: When you log out of HPE Cloud Optimizer while PKI authentication is enabled, it is recommended to close all the browser windows to successfully log in into HPE Cloud Optimizer again. If a few windows are open, the session will not be cleared and you may face issues while logging into HPE Cloud Optimizer.

The prerequisites for using HPE Cloud Optimizer with CAC are as follows:

- Server root certificate
- Client certificates
- Proxy info (if any)

To configure HPE Cloud Optimizer to use PKI authentication, follow these steps:

1. Import the trusted CA certificate to the `tomcat_trust.store` file by following these steps:

- a. Log on to HPE Cloud Optimizer as root or administrator.

- b. Run the following command:

```
<install_dir>/nonOV/jre/b/bin/keytool -import -alias
serverkey -file <CA_certificate> -keystore
"/opt/OV/nonOV/tomcat/b/ tomcat_trust.keystore"
```

In this instance, <CA_certificate> is the name (complete path) of the trusted CA certificate file on HPE Cloud Optimizer.

In the command line console, a prompt to provide password appears.

- c. Type a password, and press **Enter**.

If you do not want to set a password, press Enter without typing anything.

- d. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

A text file opens.

- e. Type the following under the **NONOV.TomcatB** section:

```
TruststoreFile=/opt/OV/nonOV/tomcat/b/tomcat_
trust.keystore
```

- f. If you typed a password in step c, type the following under **TruststoreFile**:

```
TruststorePass=<password>
```

In this instance, <password> is the password of the keystore file created in step c.

Skip this step if you did not type a password in step c.

- g. Save the file.

2. Enable checking of the validity of client certificates.

Note: If you do not perform this step, HPE Cloud Optimizer does not restrict access of users with expired or corrupt certificates, which might lead to unauthorized access.

- a. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

A text file opens.

- b. Type the following under the **NONOV.TomcatB** section:

```
OPTS_JavaOpts=-Dsun.security.ssl.allowUnsafeRenegotiation=true
-Dcom.sun.security.enableCRLDP=true
-Dcom.sun.net.ssl.checkRevocation=true -Djava.security.debug=certpath
-Dhttp.proxyHost=<Proxy_Server_IP>
-Dhttp.proxyPort=<Proxy_Server_Port>
```



```
-Dhttps.proxyHost=<Proxy_Server_IP>  
-Dhttps.proxyPort=<Proxy_Server_Port>
```

Note: Make sure that the above content is typed in a single line.

In this instance, <Proxy_Server_IP> is the IP address of the proxy server for HPE Cloud Optimizer and <Proxy_Server_Port> is the port used by the proxy server for the HTTP or HTTPS communication.

- c. Save the file.
- d. Run the following command:

```
pv enablecac true
```

HPE Cloud Optimizer is now enabled to use the PKI authentication.

Disabling Certificate-based Authentication

To disable the certificate-based authentication, run the following command at the command prompt:

```
pv enablecac false
```

Configuring LDAP

User authentication in HPE Cloud Optimizer is provided using Microsoft Active Directory and OpenLDAP. You can configure LDAP from the **Settings** page of HPE Cloud Optimizer.

Note: When you upgrade from version 2.20 to 3.00, the `PVconfig.ini` changes are retained. It is recommended to save the configuration using the interface. The LDAP configuration saved using the interface overrides the configuration saved in the `PVconfig.ini` file.

To configure LDAP for HPE Cloud Optimizer, perform the following steps:

On HPE Cloud Optimizer machine

1. From the HPE Cloud Optimizer home page, select **Settings > Integration**.
2. Scroll down to the **LDAP** section.
3. Click **Edit**.

The **LDAP Configuration** page appears.

4. Provide the following information to configure LDAP:

- LDAP Server Information

LDAP Configuration

LDAP Server Information

Host* :

Port* : ☐ SSL

Keystore Location : ⓘ

Base DN* :

User ID (Full DN) : ⓘ

User Password :

LDAP Search Properties

User Name Attribute* : ⓘ

User Search Base : ⓘ ☒ Search Subtree ⓘ

User Search Filter* : ⓘ

Group Membership* : ⓘ

Administrator Groups DN : ⓘ

User Groups DN* : ⓘ

Admin Names : ⓘ

ⓘ

*Before you save the configuration, look up for the user to validate the settings. For more information refer documentation. [Help](#)

Field	Description
Host	The fully-qualified LDAP server domain name (server.domain.com) or IP address. This is a mandatory field.
Port	The port used to connect to the LDAP server. The default port number for LDAP and LDAPS server is 389 and 636 respectively. This is a mandatory field.
SSL	If the LDAP server is configured to require LDAPS (LDAP over SSL), select the SSL checkbox.
Keystore Location	Location of the Keystore that stores the LDAP server certificate. This field is mandatory if the SSL check box is selected.

Field	Description
	<p>To get the Keystore location, perform the following steps:</p> <ol style="list-style-type: none"> Obtain the Server Certificate <p>You must add the Microsoft Active Directory server SSL certificate to the list of accepted certificates used by the HPE Cloud Optimizer server. To add the certificate, export the certificate by running the following command on the Microsoft Active Directory server:</p> <pre>certutil -ca.cert <sample.crt></pre> <p>where <code>sample.crt</code> is the name of the SSL certificate that you want to export to the HPE Cloud Optimizer server.</p> Import the Server Certificate <p>You must import the Microsoft Active Directory server certificate into the keystore for SSL enabled communication between the HPE Cloud Optimizer and Microsoft Active Directory.</p> <p>The keystore file, for example, <code>jssecacerts</code> is located in the <code><data_dir>/conf/perf</code> directory.</p> <p>Note: The <code>keytool</code> command creates the keystore file if it is not present in the directory.</p> <p>To import the server certificate, perform the following steps:</p> <ol style="list-style-type: none"> Go to the <code><data_dir>/conf/perf</code> directory. <p>Copy the exported Microsoft Active Directory server SSL certificate and paste the certificate in the directory.</p> Run the following command: <pre>/opt/0V/non0V/jre/b/bin/keytool -importcert -keystore jssecacerts -file /root/cacert.pem.</pre> <p>The <code>keytool</code> prompts you for a password. The default password is <code>changeit</code>.</p> <p>Select yes to confirm the key import when prompted with <code>Trust this Certificate?[no]: yes</code></p> Set the Keystore Location <p>Set the Keystore Location as <code><data_dir>\conf\perf\jssecacerts</code>.</p>
Base DN	The Base Distinguished Name represents the topmost level of

Field	Description
	the LDAP directory from where the LDAP search begins. For example, dc=mydomain,dc=com
User ID (Full DN)	The full Distinguished Name of the user with search permissions. This field is optional if the LDAP server does not require a User ID or Password for authentication. For example, cn=Manager,dc=mydomain,dc=com Note: In case of Microsoft AD LDS, the User DN must have "Readers Role" or "Administrators Role".
User Password	Password of the User ID. This field is optional if the LDAP server does not require a User ID or Password for authentication.

- LDAP Search Properties

Field	Description
User Name Attribute	The LDAP attribute whose value is used to log on. This is a mandatory field. For example, UID
User Search Base	The search base defines the starting point for the search in the directory tree. This value is relative to the Base DN. For example, ou=users.
Search Subtree	The Search Subtree setting controls the depth of the search under User Search Base. Select the check-box to set the scope of the search as subtree. If you want to search for a matching user in the User Search Base and all subtrees under the User Search Base, leave the Search Subtree checkbox selected. If you want to restrict the search for a matching user to only the User Search Base, excluding any subtrees, clear the Search Subtree checkbox.
User Search Filter	The pattern that is used to refine the search performed on the search base. This is a mandatory field. For example, uid={0}
Group	Attribute that contains membership information. This is a

Field	Description
Membership	mandatory field. For example, member.
Administrator Groups DN	The Distinguished Name of the Group or Organization Unit (OU) of users with administrator rights. The value is relative to the Base DN value. Multiple values should be separated with the sign. For example, cn=adminGroup1,ou=users cn=adminGroup2,ou=users.
User Groups DN	The Distinguished Name of the Group or Organization Unit (OU) of users without administrator rights. The value is relative to the Base DN value. Multiple values should be separated with the sign. This is a mandatory field. For example, cn=Group1,ou=users cn=Group2,ou=users.
Admin Names	The users that will have administrator privileges for HPE Cloud Optimizer. Admin Name is value of User Name Attribute, such as UID or cn. You can specify comma separated values. For example, admin1,admin2,admin3

5. Type the user name and click **Look Up User**. A successful search for a valid user indicates that the configuration settings are valid and HPE Cloud Optimizer can communicate with the LDAP server.
6. Type the User ID and Password of the LDAP administrator. Click **Save**. LDAP configuration saved successfully message appears.

Note: The LDAP configuration saved using this interface overrides the configuration saved in the PVconfig.ini file.

Removing the LDAP configuration

To remove the LDAP configuration for HPE Cloud Optimizer, follow these steps:

1. Log on to HPE Cloud Optimizer as a root user.
2. Run the following command at the command prompt:

```
<install_dir>/bin/ovconfchg -edit
```

A text file opens.
3. Delete the LDAP namespace.
4. Save and close the file.

Accessing HPE Cloud Optimizer User Interface

To launch the Microsoft Active Directory enabled HPE Cloud Optimizer, you can use the following URL:

http://<system>:<port>/PV or **https://<system>:<port>/PV**

If you launch the Microsoft Active Directory enabled HPE Cloud Optimizer using the URL **http://<system>:<port>/PV**, the system re-directs you to the login page using the URL **https://<system>:<port>/PV**.

Type the user name and password on the login page. For example, **PV_Admin** and **password**.

The request to validate the user name and password is sent to the Microsoft Active Directory server. The HPE Cloud Optimizer Home Page is displayed and the URL returns to **http://<system>:<port>/PV**, if the validation is successful. The protocol is redirected from HTTP to HTTPS to send the credentials to HPE Cloud Optimizer server over SSL. However, if you initiate HPE Cloud Optimizer using the URL: **https://<system>:<port>/PV**, HPE Cloud Optimizer continues to run in the secured mode.

Daily Maintenance

HPE Cloud Optimizer performs maintenance tasks on a daily basis for the following:

- To reclaim disk space for files and directories that are no longer needed
- To close any datasources that are not accessed in the last 24 hours
- To perform some basic maintenance tasks on the database

The maintenance is scheduled to run on a nightly basis at 03:00 AM local time. To configure the maintenance schedule, perform the following steps:

1. Open the *PVconfig.ini* file available in the <data_dir>/conf/perf directory.
2. Specify the hour and minutes at which you want to schedule the maintenance task, for the **DAILY_MAINTENANCE_TIME** parameter, in the 24-hour format. The default value for this parameter is 03:00.
3. Save the file.
4. Restart HPE Cloud Optimizer.

The disk space cleanup is also done for customers and users that are no longer configured. For each customer and user, there are directories that contain the customer or user graphs. The files specific to the customer logons can be located in the following directories:

- For each customer, there is a configuration directory `<data_dir>/conf/perf/VPI_CUST_<customername>`
- For a blank customer, the user directories are in `<data_dir>/conf/perf/VPI_USER_<username>`
- For users belonging to a specific customer, the user directories are in `<data_dir>/conf/perf/VPI_CUST_<customername>/VPI_USER_<username>`.

The directories along with the files that are saved are automatically removed, if the customer or user is no longer configured. To preserve these files after deleting a customer or user, you can save it to an alternate location on your local system.

Time Zones

If the time zone settings on HPE Cloud Optimizer server are different from those on the datasource, the settings on HPE Cloud Optimizer are used. The time zone settings on HPE Cloud Optimizer server are used to display data, if HPE Cloud Optimizer server and datasource are in different time zones. The default time on the HPE Cloud Optimizer server is the Coordinated Universal Time (UTC). You can change it to the time zone of your browser.

Chapter 4: Dashboard

After adding datasources, when you launch HPE Cloud Optimizer, the Dashboard view opens. The HPE Cloud Optimizer dashboard provides a quick overview of the performance and usage of entities in your environment. Following are the key features of the HPE Cloud Optimizer Dashboard:

- **At-a-glance overview**

Separate tables and graphs to quickly view the Performance, Capacity, Forecast, and Alerts. It also provides information of your environment and you can drill down for details.

- **Create your Dashboard View**

Customize the dashboard to view data specific to your role or requirement.

For example, if you are an IT Capacity Planner, you may want to see only the Capacity related information when you log in to HPE Cloud Optimizer. After logging in to HPE Cloud Optimizer for the first time, you can customize the tables and graphs that must appear when you log on to HPE Cloud Optimizer. The next time you log in, you will see only the customized data on the Dashboard page.

Click **Configure View** to reset or update the dashboard contents by selecting the components on the right side of the dashboard page.

Persona

- Virtualization Administrator
- IT Planner

Navigation

To access the feature, select **Overview**. From the drop-down, select **Dashboard**.

Using the HPE Cloud Optimizer Dashboard

The following table lists the tasks you can perform from the HPE Cloud Optimizer Dashboard:


Tasks	Description
View Treemap	Provides an overview of the performance of the entities in your environment and helps in taking corrective actions. For more information, see Use Case: Viewing Data in the Treemap


Tasks	Description
Monitor Environment Performance	<p>By analyzing the following:</p> <ul style="list-style-type: none"> • Alerts by Entity - Provides an overview of the alerts generated for the respective entities in the domain. • VM count - Provides the total number of VMs that are in Powered Off, Mostly Powered Off, and Powered On state. • VM/Host Resource Usage - Provides an overview of the CPU and memory utilization for the particular VM or Host. • Top Datastore/VM IO By Latency - Provides an aggregated value of the read and write latency consumed by the datastore or VM. • Top Datastore By FreeSpace - Provides an aggregated value of the used and free space consumed by the datastore. <p>For more information, see Monitoring Environment Performance</p>
Monitor Environment Capacity	<p>By analyzing the Capacity, you can get an overview of the total VMs available for the selected domain. You can also check the total VMs utilization and derive the allocation and usage of the remaining VMs. For more information, see Monitoring Environment Capacity</p>
View Forecast data	<p>By analyzing the Days to Capacity, you can view the number of days the VM will reach the maximum resource capacity. For more information, see Viewing Forecast Data.</p>
View Alerts	<p>Provides an overview of the critical or major alerts generated for Hosts, VMs, or Datastores. For more information, see Viewing Alerts</p>

Monitoring Environment Performance

The following table lists the possible ways to view performance data from the HPE Cloud Optimizer Dashboard:

Action	Result	Further Action, if any
View the Alerts By Entity	<p>To get an overview of the total number of alerts that are in Critical, Major, Minor, and Warning states. The alerts generated are displayed based on the Entity.</p> <p>Example</p>	<p>In the Alert messages page, you can view all the alerts corresponding to different entities and their detailed</p>

Action	Result	Further Action, if any															
	<div><div>Alerts by Entity</div><table><thead><tr><th>Entity</th><th>Critical</th><th>Major</th><th>Minor</th><th>Warning</th></tr></thead><tbody><tr><td>DATASTORE</td><td>0</td><td>6</td><td>0</td><td>0</td></tr><tr><td>GUEST</td><td>3</td><td>0</td><td>0</td><td>20</td></tr></tbody></table></div> <div><h3>Derivation</h3><p>Guest entity has 2 Critical alerts and if you want to view the cause of the alert, navigate to View All Alerts.</p></div>	Entity	Critical	Major	Minor	Warning	DATASTORE	0	6	0	0	GUEST	3	0	0	20	<p>description.</p> <p>Alternatively, you can also click the Entity to view filtered alerts based on the specific entity.</p> <p>The alerts are generated based on sensitivity. For more information, see Sensitivity</p>
Entity	Critical	Major	Minor	Warning													
DATASTORE	0	6	0	0													
GUEST	3	0	0	20													
View the VM Count	<div><p>To get a total count of the VMs in your environment that are in the Powered On, Off, and Mostly Powered Off states.</p><h3>Example</h3><div><div>VM Count</div><p>Last updated at 11/30/15 10:23</p><p>View Details</p></div><div><h3>Derivation</h3><p>Currently 140 VMs are in the Powered On state. To use the resources optimally, navigate to the PowerState for VM page.</p></div></div>	<p>Use the VM count page to get detailed information on the performance of the individual VMs.</p> <p>Based on the analysis, you can reclaim the possible VMs that are mostly Powered Off and Powered On for longer periods, thereby improving the performance of your environment.</p>															
View the VM/Host Resource Usage	<div><p>To get an overview of the maximum CPU and memory utilization for a particular VM or Host.</p><h3>Example</h3></div>	<p>Based on the analysis, you can allocate the resources to improve the performance of your environment.</p>															


Action	Result	Further Action, if any
	 <p>Derivation</p> <p>The scattered graph displays the CPU and memory utilization for the particular VM. If you find the resource utilization high or low, you can plan the allocation of the resources to suit your environment needs.</p>	
View the Top Datastore/VM IO By Latency	To get an overview of the read/write latency in milliseconds on the datastore or VM. Lists the top five most utilized resources.	
View the Top Datastore By Freespace	To get an overview of the used or free space in GB on the datastore. Lists the top five most utilized resources.	

Monitoring Environment Capacity

The following section details information on using the HPE Cloud Optimizer Dashboard to view the capacity and usage data of entities. You can view a summarized or detailed analysis of your existing infrastructure that helps you improve the efficiency of resources by optimal sizing and placement of the VMs.

The following table lists the possible ways to view Capacity and Usage data from the HPE Cloud Optimizer Dashboard.

Action	Result	Further Action, if any
View the Capacity	To get an overview of the total remaining capacity of VMs for the monitored domain.	To analyze how the environment would perform

Action	Result	Further Action, if any
	<p>If you see that the remaining capacity for the VMs is less, you can take appropriate actions from the data derived.</p> <p>Example</p>  <p>Capacity</p> <p>Deployed Remaining</p> <p>vSphere OneView KVM</p> <p>0 50 100 150 200 250 300 347</p> <p>Last updated at 11/30/15 10:29</p> <p>Unified Enclosure View Model Scenario Capacity Details</p> <p>Derivation</p> <p>There are 228 VMs deployed and remaining capacity of 20 VMs for a vSphere domain. To view the CPU, memory, and storage utilization of the VMs navigate to the Capacity Details page and to allocate the resources navigate to Model Scenario page.</p>	after allocating or removing the resources.

Viewing Forecast Data


The following table lists the possible ways to view Forecast data from the HPE Cloud Optimizer Dashboard:

Action	Result	Further Action, if any																
View the Days to Capacity	<p>To get an overview of the entities reaching capacity in the next few days.</p> <p>You can analyze the table and take required actions based on the data.</p> <p>Example</p> <div><div>Days to capacity</div><table><tr><th>Entity</th><th>Entity Type</th><th>Days</th><th>Constraint</th></tr><tr><td>Mordor</td><td>Datacenter</td><td>0</td><td>Memory</td></tr><tr><td>Mordor Cluster</td><td>Cluster</td><td>0</td><td>Memory</td></tr><tr><td>Business</td><td>Cluster</td><td>0</td><td>Storage</td></tr></table></div> <p>Derivations</p>	Entity	Entity Type	Days	Constraint	Mordor	Datacenter	0	Memory	Mordor Cluster	Cluster	0	Memory	Business	Cluster	0	Storage	With the forecast data of the entities, you can plan the resource allocation effectively.
Entity	Entity Type	Days	Constraint															
Mordor	Datacenter	0	Memory															
Mordor Cluster	Cluster	0	Memory															
Business	Cluster	0	Storage															

Action	Result	Further Action, if any
	2 datastores and 1 cluster reached or reaching capacity in the next 11 days. Based on the constraint you can view the entity that has reached the maximum capacity. Navigate to the Aggregate Forecast page to get an overview of the CPU, Memory, and Storage utilization and make necessary changes to the entity reaching capacity.	

Viewing Alerts

The following table lists the possible ways to view alerts from the HPE Cloud Optimizer Dashboard:

Action	Result	Further Action, if any
View the top pane listing the Critical and Major alerts corresponding to entities.	<p>To view if there are any alarming issues in the environment that need immediate attention.</p> <p>Example</p>  <p>The screenshot shows a row of four entity cards: '17 VMs', '361 VMs', '44 Subsystems', and '44 Subsystems'. Each card displays a 'Critical' count (red square) and a 'Major' count (orange square). For '17 VMs', Critical is 0 and Major is 0. For '361 VMs', Critical is 3 and Major is 0. For '44 Subsystems', Critical is 0 and Major is 0. For '44 Subsystems', Critical is 0 and Major is 0.</p> <p>Derivations</p> <p>0 Critical and 7 Major alerts corresponding to the VMs in the environment. Navigate to the Alert messages to view detailed information about the alert or Workbench to troubleshoot the problem.</p>	To get detailed information about the alert. You can also troubleshoot the problem or draw graphs for the particular metrics.

Customizing the Dashboard View

After logging into HPE Cloud Optimizer, you can customize the Dashboard View to show only those cards which are relevant to your area of expertise or profile. To customize the view, perform the following steps:

1. Click **Configure View** available at the right side of the page. The **Select Dashboard Components** pane appears.

Select Dashboard Components

- ☐ Alerts by Entity
- ☐ Capacity
- ☐ Days to capacity
- ☒ VM Count
- ☐ Top 5 Datastores By IO Latency
- ☒ Top 5 VMs By IO Latency
- ☐ Top 5 Datastores By FreeSpace
- ☐ VM Resource Usage
- ☐ Host Resource Usage
- ☐ DSDeviceLatencyComponent

Done

2. Clear the check boxes corresponding to the cards you do not want to view.
3. Click **Done**.
4. Refresh or reload the page to view the Dashboard with only the selected cards.

Note: Alternatively, you can also close a card to remove it from the Dashboard view. The closed card is not available when you log in the next time.

Chapter 5: Treemap

The Treemap is a visual representation of the entities in your environment. Treemap displays the hierarchical data as a set of nested rectangles. The color and size dimensions used for the leaf nodes often get visually correlated to spot interesting patterns.

You can use the Treemap for a quick analysis of the performance of the entities and take corrective actions. You get a domain specific view of your environment. You can monitor various domains such as vSphere, KVM/Xen, Hyper-V, Business Groups, OneView, and Physical servers. To view data in the Treemap, add the datasources to be monitored and then wait for the first data collection to complete.

You can do the following tasks from the Treemap.

- Identify whether the problem can be isolated to a particular type.
- Identify patterns of problems across elements of similar type.
- Drill down from the Treemap to find the reason and status of the resource allocation and usage and to take further actions. See, [Real Time Guest OS Drilldown for Advance Troubleshooting](#).
- Displays the resource utilization of the resource type you selected in the Resources pane.

For more information on adding a datasource, see [Adding a Datasource](#).

Navigation

To access the feature, click **Overview> Treemap**.

Treemap View

The following table lists the view details of the Treemap Overview page and the corresponding functions:

User Interface Feature	Description
Resources Pane	Displays the resources based on which you can view data in the Treemap. For more information, see Resources Pane .
Options Pane	Contains the fields and filter based on which you can view data in the Treemap. It also contains the Micro charts feature.

Note: The minimum screen resolution required for viewing the HPE Cloud Optimizer Treemap is 1280x768.

Using Treemap View

The Treemap view provides a visual representation of the performance of resources monitored.

You can perform the following tasks on the Treemap:

Use	Description
Viewing data from the Treemap	You can view the data specific to a resource by right-clicking the entity.
Viewing resource information	You can view information corresponding to the specific resources in a pop-up window.
Drill down	Navigate to display the resource utilization specific to the resources further down in the hierarchy.
Viewing Micro charts	Quickly analyze and compare the performance or utilization of your monitored resources. To view the detailed graphs for different attributes, you can use the Workbench page.
Searching the Resources	There are various ways that are available to search the resources in the environment.

You can use the different options in the **Options** pane to filter your view in the Treemap:

- Group the resources displayed
- Set attributes based on which the Treemap is colored
- Filter data shown in the Treemap
- View data up to the past five summarization intervals
- List up to five most and least utilized resources
- Highlight up to five most and least utilized resources

Viewing data from the Treemap

From the Treemap, right-click the specific resource to view the data. The context-

sensitive menu includes the following options:

- **Utilization Trends:** Displays the default graphs for the selected resource.
- **Show Status:** Displays the Status Report for the selected resource by default. You can select and view other corresponding reports also.
- **Real Time Guest OS Drilldown:** For more information, see ["Real Time Guest OS Drilldown for Advance Troubleshooting "](#) on page 238.
- **Capacity Usage Trends:** Displays the overview of the resources optimization.
- **Sizing Recommendation:** Displays the sizing data for the resource.
- **Forecast:** Displays the forecast data for the resource.
- **Drill Down:** Displays data corresponding to the hosts and VMs for the selected resource.

Note:

- The **Drill Down** option is available only when you are viewing the Treemap data corresponding to certain resources in your datasource.
- If you want to view the VMs corresponding to a host, right-click the host and select **Show Status**. The Status Report opens and you can view the VMs corresponding to the host.
- The Treemap does not show those resources for which no data is available.

- **Highlight Element:** Highlights the selected element(s) in the Treemap. To highlight multiple elements, right-click the element and select the option.
- **Highlight Group:** Highlights the group in which the selected element(s) is grouped. You can use this option only when you use the **Group By** feature.

Note: If a resource does not belong to any group, then it will be listed under a custom group called **Ungrouped**.

- **Remove Highlights:** Removes the highlighting for the element(s) or group.
- **Launch Workbench:** Launches the Workbench (Performance Graphing) page.
- **Launch Workbench (Highlights):** Launches the Workbench page for the highlighted element(s) or group.

Resource Information in Treemap

When you move the pointer over resources in the Treemap, you can view information corresponding to the specific resources. This window displays the following data:

- Name of the resource
- Value corresponding to the attribute selected in the **Color by** field in the **Options** pane.

- Value corresponding to the **Size by** attribute of the resource, specified in the **Options** pane.

Related Topics:

- ["Use Case: Viewing Data in the Treemap" on page 114](#)

Resources Pane

The Resources Pane on the top of the HPE Cloud Optimizer home page displays the resources for which you can view the utilization. The pane also displays the total number of resources available for each type and the corresponding micro charts. The attributes for each of the resources, based on which you want to view the Treemap can also be selected by clicking the corresponding icons.

Group by and Color by

After selecting a resource, you can further view specific data in the Treemap by selecting the attributes available in the **Group by** and **Color by** fields. Alternatively, you can also select the icons corresponding to each resource in the Resources pane.

The following table lists the **Group by** and **Color by** attributes corresponding to the resources for the different datasources.

Domain	Resource	Group By	Color By
vSphere	Datcenters	None	<ul style="list-style-type: none">• VMs Offline• CPU Usage• Memory Usage
	Cluster	Data Center	<ul style="list-style-type: none">• CPU Usage• Memory Usage
	ESX/ESXi Hosts	<ul style="list-style-type: none">• Data Center• Cluster	<ul style="list-style-type: none">• CPU Usage• Memory Usage
	Resource Pools	<ul style="list-style-type: none">• Data Center• Cluster	<ul style="list-style-type: none">• CPU Usage• Memory Usage
	VMs (Compute)	<ul style="list-style-type: none">• Data Center• Cluster• Resource Pool	<ul style="list-style-type: none">• CPU Usage• Memory Usage• CPU Ready

Domain	Resource	Group By	Color By
	Datastores	Data Center	Usage
	VMs (Storage)	<ul style="list-style-type: none"> Data Center Storage 	Latency
Hyper-V	Host Groups	None	<ul style="list-style-type: none"> CPU Usage Memory Usage
	Host Clusters	Host Group	
	Hosts	<ul style="list-style-type: none"> Host Group Host Cluster 	
	VMs	<ul style="list-style-type: none"> Host Group Host Cluster Host 	
KVM	Hosts	None	CPU Usage
	Virtual Machines	Hosts	
Xen	Hosts	None	CPU Usage
	Virtual Machines	Hosts	
OpenStack	Cloud	None	<ul style="list-style-type: none"> CPU Usage Memory Usage
	Tenants	Cloud	<ul style="list-style-type: none"> CPU Usage Disk Usage
	Hypervisors	None Cloud	<ul style="list-style-type: none"> CPU Usage Memory Usage Disk Usage
	VMs (Compute)	None Tenant Cloud	<ul style="list-style-type: none"> CPU Usage Memory Usage

Domain	Resource	Group By	Color By
OneView	Enclosure	None	Power Usage
	ServerHardware	Enclosure	<ul style="list-style-type: none">• CPU Usage• Power Usage
Physical Server	Server(s)	None	CPU Usage
AWS	Account(s)	None	CPU usage
	Instance(s)	None	CPU usage

Note: In the HPE Cloud Optimizer Treemap, after you drill down from a resource, change the color by value for a resource using the Color By drop-down list in the Options pane. If you change the color by value using the resource name in the Resources pane, the drill down and drill up features won't work in the Treemap.

For example, in the vSphere domain:

1. Drill down from the Cluster level to the host level with Color By value as **CPU Usage**.
2. Now, if you change the color by value to **Mem Usage** using the resources name in the Resources pane, you cannot drill down or drill up to any other resources.

Hence, to drill down to the VMs view from the hosts view, use the Color By drop-down list in the Options pane to change the color by value.





Using the Options Pane


The Options pane provides you the fields and values based on which you can view and interpret data in the Treemap. You can use the different options in the **Options** pane to filter your view in the Treemap:



- Group the resources displayed
- Set attributes based on which the Treemap is colored
- Filter data shown in the Treemap
- View data up to the past five summarization intervals
- List up to five most and least utilized resources
- Highlight up to five most and least utilized resources

Options Pane Details

The following table explains the fields available in the Options pane.

Field	Description
Search	Use the text box to search for resources within the Treemap. The Treemap view shows the searched resources highlighted. For more information, see Searching Resources .
Remove Highlights ()	Clears the text box corresponding to the Search feature and also clears the Treemap highlighting resulting from the previous search.
Export as PDF ()	Click to export all the data in the page, including the Treemap, as a PDF document.
Help ()	Click to open the context-sensitive help for the page.
Domain	Select the domain based on which you want to view data in the Treemap.
Group by	Groups data in the Treemap based on the value selected from the list. The values in the drop-down list change based on the domain and resource selected on the Resources pane. For more information, see Resources Pane .
Data as of	Displays the date and time corresponding to the data collected and displayed in the Treemap. Use the slider to set the required time and the Treemap displays the corresponding data. You can view data up to the past 5 summarization intervals.
Summarized Every	Displays the time interval configured for collecting data. The time interval varies depending on the datasource domain monitored.
 (Information)	<p>Displays additional information about your resources and other attributes. To view the information, move the pointer over the icon.</p> <p>In the Resources pane, moving the pointer over the icon shows you the status and count of the resource selected. To view the icon for the resource pools, click Resource Pools on the pane.</p>
Size by	The attribute based on which the sizes of the boxes in the Treemap are drawn. The attribute changes based on the value selected in the Color by field.

Field	Description
Color by	Colors data in the Treemap based on the attribute selected from the list. The attributes in the list vary based on the resource selected on the Resource pane.
Filter	Filters the displayed data based on the range set between the end points on the slider. For more information, see Filtering Data .
Spectrum Transition Value	<p>Spectrum Transition Value is the value set on the Green-Yellow-Red color spectrum which defines the optimum performance area for a resource.</p> <p>The area surrounding the Spectrum Transition Value is Yellow. This indicates that any resource having performance parameters within the Yellow range is performing optimally. The resources having optimal performance parameters are colored yellow in the Treemap.</p>  <p>For example, in the above figure, the Spectrum Transition Value is set to 75. Hence, when you view VMs in vSphere by CPU Usage, then all the VMs having CPU utilization in the Yellow range will be optimally used VMs.</p>
Hide Labels	Select this check box to hide the labels of the resources displayed in the Treemap. To view the labels, clear the check box.
Top	Select the radio button to view the top 5 utilized resources, based on the Color by attribute, in the current Treemap view. The pane shows the resource names and the corresponding graphs.
Bottom	Select the radio button to view the least 5 utilized resources, based on


Field	Description
	the Color by attribute, in the current Treemap view. The pane shows the resource names and the corresponding graphs.
Highlight	Highlights the Top or Bottom listed resources in the Treemap. This helps you to locate the resource easily in the Treemap.
Collapse Expand	Click  in the top right corner of the Options pane to collapse the Options pane and make more space for the treemap. Click  (Expand) to expand the Options Pane again.

Drill Down to Resources

The Drill Down feature helps navigate to further available levels and view the resource utilization. When you select **Drill Down** from the context-sensitive menu in the Treemap or double-click a resource, the Treemap refreshes to display the resource utilization specific to the resources further down in the hierarchy. The Drill Down option is available only when you are viewing specific resources in the Treemap.

For example, for a monitored VMware vCenter Server, from the Treemap view for a cluster or resource pool, if you click **Drill Down**, the Treemap refreshes to display the resource utilization of the corresponding vSphere hosts. To view the Treemap data corresponding to the VMs in the particular cluster or resource pool, you can click **VM (Compute)** in the Resources pane. Alternatively, you can also switch the views by selecting the required value from the drop-down list that appears at the top of the Treemap once you drill down.

From the Treemap view for a data store, if you click **Drill Down**, the Treemap refreshes to display the resource utilization of the VMs corresponding to the selected data store.

Click  (Drill Up) to go back to the cluster, resource pool, or data store view. Alternatively, you can also right-click the Treemap, and select **Drill Up**.

Note: Double-clicking a resource shows you the corresponding drilled-down data. If there are further levels to be drilled-down, double-clicking again shows that data. When you reach the last level of the drill-down and again double-click a resource, it drills up.

Nested Resource Pools

If you double-click a resource pool having a nested resource pool in it, the Treemap refreshes to display the nested resource pool. Double-clicking the nested resource pool,

further displays the corresponding VMs. To view the VMs corresponding to the first level of drill-down, select **VMs** from the drop-down list that appears at the top of the Treemap once you drill down.

Cross launch to actual domain from Treemap

While using the Treemap to view the virtualized entities in the **OpenStack** domain, you can cross launch into the actual domain of the Hypervisor and the VM. The OpenStack VMs and Hypervisors can belong to various virtualization domains. You can view the detailed set of metrics from the actual domain (VMware or KVM), view forecast, capacity summary, and sizing recommendations for the VMs.

To cross launch to the actual domain of the VM and Hypervisor from the OpenStack view of the Treemap:

Right-click the Openstack VM or Hypervisor, and select Launch Workbench, Forecast, or Capacity Usage Trends.

Note: You can view the metrics for the OpenStack domain by accessing the VM or Hypervisor directly from the Workbench.

To view the OpenStack metrics, perform the following steps:

1. From the HPE Cloud Optimizer user interface, select **Performance > Workbench**.

The Workbench appears.

2. Select the OpenStack entity from the left pane.

The list of metrics appear in the Metrics pane.

Searching Resources

You can search for a resource available in the Treemap, using the **Search** box in the **Options** pane. Type the name of the resource or a regular search expression in the text box. The resources matching the search criteria are highlighted in the Treemap.

Note: The search feature does not highlight resources in the Treemap if you have already selected the highlight option to view the top or bottom resources in the Treemap.

You can search for a resource by typing the first alphabet of the resource name or any of the characters of the node name. The text you type in the text box is not case sensitive.

Following are a few examples:

- To search for the resources having *Virtual* in the resource name, type `virtual` in the text box.
- To search for all resources having the resource name starting with *Virtual*, type `^virtual` in the text box.
- To search for all resources having the words *Virtual* and *app* in the resource name, type `virtual.*app` in the text box.

You can use the regular search expressions to search for resources in the Treemap.

The Treemap appears grayed out, if the resource you are searching for is not available in the Treemap.

The Search state remains saved when you click other resources in the Resources pane. For example, if you use a search criteria while clusters is selected, and then click datastores, the resources corresponding to the search appear highlighted in the Treemap view for datastores.

To clear the search criteria typed in the text box, click **(Remove Highlights)**. The highlighting resulting from the previous search is removed from the Treemap.

Filtering Data

Filtering enables you to filter the already shown data in the Treemap. You can set a range, between which you want to display the data in the Treemap.

To filter data in the Treemap, follow these steps:

1. Select the required resource from the Resources pane.
2. Select the value based on which the data must be grouped in the tree map in the **Group By** field.
3. Select the value based on which the data must be displayed in the **Color By** field.
4. Click and drag the end points of the Filter slider and set it to a required range.

The Treemap refreshes to display the data based on the set filter.

Micro Charts

The micro charts in HPE Cloud Optimizer help the user to quickly view trends of the resources being monitored. These micro charts are graphs and correspond to the selected resources and set attributes. Using these micro charts you can quickly analyze

and compare the performance or utilization of your monitored resources. To view the detailed graphs for different attributes, you can use the Workbench page.

Viewing Micro Charts

The micro charts always appear in accordance with the top or bottom resources, based on your selections in the **Options** pane. For every Treemap view, the HPE Cloud Optimizer interface lists the top or bottom resources. When you select **Top** or **Bottom**, HPE Cloud Optimizer lists the corresponding resources from the Treemap view and the associated graph.

The value selected for the **Color by** field in the Options pane is the attribute based on which the graph is drawn. For more information on the values in the **Color by** field, see ["Group by and Color by" on page 106](#)

Use Case: Viewing Data in the Treemap

This section shows, how you can use the features provided by the Treemap to view the required data.

Scenario


John is a VMware Administrator in an organization. He wants to view the top three datastores that are utilizing 80% to 90% of their storage space and view the storage virtual machines associated with these three datastores.

The tasks that John must perform as an Administrator to accomplish his objectives are as follows:

1. Log on to the HPE Cloud Optimizer home page.
2. In the Resources Pane, click **Data Stores**. The Treemap refreshes to show all the datastores corresponding to the added VMware vCenter servers.


The **Usage** attribute decides the coloring in the Treemap.

3. Move the end point values on the filter slider to 80 and 90.



The Treemap refreshes to display the datastores that have the storage usage between 80% and 90%.

4. Select the **Top** option. The names and graphs corresponding to the top five datastores that have storage usage in 80% to 90% range are listed.

5. Select **Highlight**. The Treemap highlights the top five utilized datastores in the selected range.
6. Move the pointer over the highlighted datastores to view the exact percentage utilization of the top three datastores.
7. To view the storage VM corresponding to each data store, right-click the data store and select **Drill Down**. The Treemap refreshes to display the storage VMs associated with the selected data store.
8. To go back to the data store view, click  (Drill Up) icon.

Chapter 6: Performance

HPE Cloud Optimizer collects data from the added datasources and provides the performance-related information of the resources. You can use this performance data to perform the following actions:

- Identify and troubleshoot problems in your environment
- Plan and allocate your resources better
- Get an overview of resource utilization
- Compare performance of resources across domains and time periods

The HPE Cloud Optimizer features using which you can view the performance data are as follows:

- **Treemap** - Provides you a visual representation on the performance of the resources.
- **Performance Graphs** - Provides graphs drawn based on the data collected from the datasources.
- **Reports** - Provides data in the form of reports.
- **Alerts** - View the alert messages in the environment.

Navigation

To access the feature, from the HPE Cloud Optimizer home page, select **Performance**.

Performance Graphs

Navigation

To access the feature, from the HPE Cloud Optimizer home page, select **Performance > Workbench**.

The Workbench page helps you to draw graphs for resources you are monitoring.

The preformatted and custom graphs help you perform the following:

- Evaluate the performance of monitored entities (such as applications, system, and network)
- Analyze usage trends
- Correlate usage
- Compare performance of different resources

You can also compare multiple instances of a resource on one or more Configuration Items (CIs). CIs are the objects that you are monitoring in your virtualized environment.

Graphs Overview

A graph family tree organizes the Performance Graphing and consists of the following:







- **Family**
Refers to the group that organizes the graphs.
- **Category**
Refers to the sub-groups of graphs that are logically grouped within the family.
- **Name**
Uniquely identifies a graph definition.







Default Graphs

A default graph in a graph family contains metrics to measure the performance of any resource or application. You can map graph families or categories to a CI. When you launch a graph for a CI that has a graph family or category associated, the system draws the default graph from the graph family or category. A graph family can have one or more default graphs. If a graph family does not have a default graph, the system selects the first graph in the family or category.

Icon Glossary

The following table lists the icons used to represent resources in the **Configuration Items** tree in the Workbench page.

Icon	Name of Resource
	Datacenter
	Cluster
	Guest and BYVM_Storage
	Resource Pool
	Host (VMware vCenter, HyperV, KVM, Xen, and OpenStack)
	Host in maintenance mode

Icon	Name of Resource
	Storage
	Host Group
	Cloud (OpenStack)
	Tenants (OpenStack)
	User
	Application





Note: If an icon appears grayed-out, it indicates that the resource is in the Powered Off/Shutdown/Shutoff/Stopped state.


Workbench Page

The Workbench page provides the Performance Graphing functionality for HPE Cloud Optimizer.

Configuration Items (CI)



The Configuration Items tree displays the list of CIs. You can select the required CI for which you want to draw the graph. The following table lists the functions available from the Configuration Items tree.

Icon/Field	Description
	Click to collapse the Configuration Items pane.
	Click to expand the Configuration Items pane. This button is available only when you collapse the pane.
Refresh 	Click to refresh the CIs available in the Configurations Items tree.
Filter By 	Click to filter the CIs in the Configuration Items tree. The drop-down list contains values that are based on the type of CIs monitored.

Icon/Field	Description
	Click (Remove Filter) to remove a set filter.
Search	Type a value in this field to search for resources available in the Configuration Items tree. The Search field only appears if you have filtered the CIs.
Clear ()	Clears the text box corresponding to the Search feature and also clears the Configuration Items tree highlighting resulting from the previous search.

Favorites

The Favorites pane lists the graphs you have saved as a favorite to be accessed later. The following table lists the functions available in the Favorites pane.

Icon	Description
Refresh 	Click to refresh the list of favorites in the pane.
Delete Favorites ()	Click to delete the saved favorite graphs.


This option appears only when there are saved favorites in the pane. For more information on adding favorites, see ["Saving as Favorite" on page 129](#).

Performance Pane

The Performance pane enables you to draw and view performance graphs for a selected CI. You can draw a pre-defined graph for a CI or draw a graph from the list of available metrics.

The following tables lists the tabs available in the Performance pane.

Tab	Description
Metrics	Displays the following information: <ul style="list-style-type: none">• Metric Classes - lists the metric classes available. This data is not displayed, if there is only one metric class.• Instances - lists the instances available for the selected multi- instance metric class.• Metrics - lists the metrics available for the selected metric class

Tab	Description
	You can use the Refresh () icon available in the Metrics tab to refresh the available metric classes, instances (if any), and metrics.
Graphs	Displays a list of graph templates based on the CI selected from the Configuration Items pane.
Reports	Displays the list of reports for the selected CI. For more information, see "Reports Overview" on page 131




Options for Drawn Graphs

This section provides information to interpret the contents of graphs and use the available features and functionality to change the content.

Graph Selection Icons



The **Graphs** tab displays a list of graphs, graph families, and categories associated with the selected CI. By default, the tab selects default graphs in the graph families or categories and the CIs used.




The following table lists the functions available in the Graphs tab.

Icon	Description
 (Draw Graphs)	Draws graphs for the selected CI.
 (Clear Selections)	Clears the selections in the tab.
 (Refresh)	Refreshes the Graphs list.

Drawn Graphs Window

The following table lists the elements available in the title bar of the Performance Pane.

Icon	Description
Collapse ()	Click to collapse the Graphs , Metrics , and Reports tabs.
Expand ()	Click to show the Graphs , Metrics , and Reports tabs. This icon appears after you collapse the tabs.

Icon	Description
Options ()	Displays the menu options for the drawn graphs. For more information, see the "Options Menu" below .
Export as PDF ()	Click to export all the drawn graphs into a PDF document. This icon appears after you draw graphs on the Performance pane.
Save as Favorite ()	Click to save the graphs drawn as a Favorite and access them later. This icon appears after you draw graphs on the Performance pane.



Options Menu

The following table lists the options available in the **Options** menu on the title bar of the Performance Pane.

Option	Description
Tooltips	Select this option if you want to enable tooltips for the drawn graphs. When selected, moving the pointer on the graph area of a drawn graph opens a text box displaying the actual value of the data point and the time interval of the selected data. No pop-up window is displayed when this option is disabled.
Date Range Panel	Select this option to open the Date Range Panel. Using the Date Range Panel, you can quickly change the duration for which the graph is drawn. Use this option to drag a metric and draw a reference graph. The reference graph helps you understand the correlation between different metrics. You can also customize the graph to view granular data for a selected unit of time.
Navigation Panel	Select this option to display the navigation panel in the graph window. This option is available only for graphs that display near-real time data.
Close All Graphs	Select this option to close all open graph windows simultaneously.

Graph Window Options

The following table lists the options available in the Graph window.

Option/Icon	Description
Metric Legend	Click the Metric Legend icons to show or hide the corresponding graphs of the metric. To remove a metric from a graph, right-click the Metric Legend icon and select Remove . HPE Cloud Optimizer removes the corresponding graph and the legend from the graph window.
(Auto Refresh On/Off) 	If you enable the Auto Refresh option, the data in graphs is refreshed at a specific interval.
(Prev/Next)	Performance Graphing provides the Prev and Next buttons to navigate to adjacent time intervals.
Prepend/Append 	Performance Graphing provides the Append and Prepend buttons to append or prepend data for adjoining time intervals.
Options > View as Table	Select this option to view data in a tabular format.
Options > Export	Select this option to export graphs to formats such as .tsv, .csv, .xls, and .xml. For more information, see Graph Export Dialog Box .
Options > Navigate	Select this option to select the start and end date and time from the Time Settings dialog box. The graph refreshes to display the data for the selected period. If you select the Apply to all graphs check box, data in all the graphs refresh.

Zoom In Data in Drawn Graph

After you draw a graph, you can zoom-in to view data points for a smaller time period and zoom-out to reset and view the original graph. Zooming-in or zooming-out enables you to adjust the summarization levels.








- To zoom in, click and drag the pointer on the graph from left to right.
- To zoom out, click and drag the pointer on the graph from right to left. When you zoom out a graph, the graph resets to the previous state.

You can zoom in to multiple levels. Every zoom-out reverts the graph to the state before the zoom-in operation.

Table Graph Window

A table graph helps you to view granular data in numeric format. This section describes in detail the information displayed in the Table Graph window. To view the graph as a table, draw a graph and select **Options > View as Table** from the graph window.

The Table Graph window displays the following elements:

Icon/Option	Description
Lock Column	<p>Select this option to lock one or more columns in the table graph. This option helps you to view the locked columns as you navigate across columns using the horizontal scroll bar.</p> <p>When you select the Lock Column check box, by default, the system locks the first column for viewing. To add another column, drag an unlocked column before the thick column border that separates locked and unlocked columns. The thick column border indicates the range of columns that are locked.</p>
 (Table Highlights)	Click this icon to open the Table Highlights dialog box. You can set the attributes based on which rows in the table must be highlighted. For more information, see "Using Table Highlights" on the next page .
 (Table Filters)	Click this icon to open the Table Filters dialog box. You can set the attributes based on which rows in the table must appear. For more information, see "Using Table Filters" on page 125
 (Export)	Use this option to export the data from tabular format to .csv, .tsv, Excel, or .xml formats.
 (Auto Refresh On/Off)	If you enable the Auto Refresh option, the data refreshes at a specific interval.
 (Prev/Next)	Use the Prev and Next buttons to navigate within the table.
  (Sort Metric Columns)	Use this option to sort the data in the metric columns in the ascending or descending order. For more information, see Specifying the Order of Metric Columns below.

Specifying the Order of Metric Columns

You can specify the order of the metric columns, based on which data in the columns must be sorted. To specify the order, perform the following steps:

1. Select the metric based on which the data must be sorted, in the Table graph window.
The value **1** appears corresponding to the metric specifying the order.
2. Move the cursor to the next metric that must be considered while sorting data and click on the ▼ or ▲ (**Sort Metric Columns**) icon to set the order.

Note: This option appears only when there are two or more metrics.


3. Repeat step 2 till you sort all the required metrics in order.
The number corresponding to each metric represents the order specified.

Note: To reset the order to **1**, click any metric name. You can specify the order again.

4. Click ▼ or ▲ (**Sort Metric Columns**) icon corresponding to the metric that is assigned the order **1** to view the data in the ascending or descending order.

Using Table Highlights

You can highlight cells in the table by specifying conditions based on the metric value. To highlight cells, perform the following steps:

1. Click () **Table Highlights** icon. The Table Highlights dialog box opens.
2. Select the required value based on which the cells must be highlighted, from the drop down list.
3. Select the required comparison symbol from the list of symbols available. For more information on the comparison symbols, see "[Comparison Symbols](#)" below
4. Type the value to be compared in the text box.
5. Select a color from the color palette to highlight the cells.
6. Click **Add**. The highlight condition is added to the Conditions table.
To remove any added condition, select the condition and click **Remove**.
To remove all the added conditions, click **Remove All**.
7. Click **OK**. The Table Highlights dialog box closes.
The values matching the condition set is highlighted in the table.


Comparison Symbols

The following table lists the comparison symbols available in the Table Highlights dialog box.

Comparison Symbol	Description
<=	Less or Equal
>=	Greater or Equal
!=	Not Equal
!~	Not Like (text with leading or trailing ".*" expressions)
=	Equal
~	Like (text with leading or trailing ".*" expressions)
<	Less
>	Greater

Using Table Filters


You can use the Table Filters in the table to filter and view data in the table. To filter rows, perform the following steps:

1. Click () **Table Filters** icon. The Table Filters dialog box appears.
2. Select the required value based on which the cells must be filtered, from the drop down list.
3. Select the required comparison symbol from the list of symbols available. For more information on the comparison symbols, see ["Comparison Symbols" on the previous page](#)
4. Type the value to be compared in the text box.
5. Select either **AND** or **OR** to apply multiple conditions.
 - **AND** - filters the rows that satisfy the first condition you set and the row that satisfies the following conditions as well.
 - **OR** - filters the rows that satisfy either the preceding or the following conditions.
6. Click **Add**. The filter condition is added to the Conditions table.
7. Click **OK**. The Table Filters dialog box closes.

The table displays only the rows matching the set filter conditions.

Updating Filter Conditions

To update a filter condition in the Table Filters dialog box, perform the following steps:

1. Click () **Table Filters** icon. The Table Filters dialog box opens.
2. Select the filter condition to update from the list of conditions in the table. The filter values are updated in the preceding fields.
3. Update the values as required.
4. Click **Update**. The filter condition is updated in the table.
To remove any added condition, select the condition and click **Remove**.
To remove all the added conditions, click **Remove All**.

Graph Export Dialog Box

You can export drawn graphs to one of the following formats: .tsv, .csv, .xls, and .xml. To export graphs, draw a graph and select **Options > Export** from the graph window. The Export from a Graph window opens.

The Graph Export dialog box includes the following elements:

UI Elements	Description
Cancel	Click Cancel to close the dialog box and return to the drawn graph page.
OK	Click OK to export the data in the selected format.
Type	Select the format in which you want to export the graph.


Note:

The exported data has the same time as that of the HPE Cloud Optimizer server.

Date Range Panel

The Date Range Panel option helps you to view data for specific periods of time.





To access the Date Range Panel:



- Select **Options > Date Range Panel** from the Performance Pane.
- You can also view the panel by clicking  (**Show Date Range Panel**) in the Performance pane.

The following table lists all the options available to customize a graph.

Note: When you access the Date Range Panel, the panel displays the Start Time and End Time. This denotes the start and end time on the date range slider and not the start and end time of the data in the graphs. As you move the slider, the values

for the Start Time and End Time are updated.


Icon/Option	Description
Show Date Range Panel ()	Click to view the Date Range Panel.
Last <Unit of time>	<p>Select a unit of time from the list to view granular data for that period of time. The options available are as follows:</p> <ul style="list-style-type: none"> • Hour • Day • Week • Month <p>For example, if you select the value Hour, Performance Graphing draws the graphs for the last one hour.</p>
Range	<p>Select the time range from the list. The options available are as follows:</p> <ul style="list-style-type: none"> • Months in Year • Weeks in Month • Days in Week • Hours in Day <p>By default, a Range value is based on the selected unit of time.</p> <p>For example, if you select last hour from the list, the Range list displays hours in a day.</p>
 (Choose a time frame)	Select a start date and end date from the Time Settings dialog box to view data for that period. The start date and end date values range between the earliest and the latest time stamps available in the data source.
 (Previous)	<p>Select Previous to view the data for the adjacent unit of time. For example, if you select last hour from the list, you can use the Previous Hour option to view data for that period.</p> <p>In the same way, you can view adjacent data for any unit of time.</p>
 (Next)	Select Next to view the data for the adjacent unit of time. For example, if you select last hour from the list you can use the Next

Icon/Option	Description
	Hour option to view data for that period. In the same way, you can view adjacent data for any unit of time.
 Adjust to a unit of time	Use this option to move the slider to a unit of time you select from the Unit of time list. For example, if you select an hour from the list, you can use adjust to an hour and the slider moves to show data for an hour
Apply to all graphs	Select this option to apply the changes made within the date range panel to all the drawn graphs. If you do not select this option, the changes apply only to the selected graph and the selected graph appears highlighted.
Hide Date Range Panel ()	Click this icon to hide the Date Range Panel.

Drawing Graphs

You can launch graphs from predefined graph templates or create custom graphs by selecting the required metrics for a selected CI. You can draw a graph from the Workbench page in one of the following ways.

To draw graphs from Graph templates:

1. From the **Configuration Items** tree, select the CI for which you want to draw graphs.
2. From the Performance pane, select **Graphs** tab. This tab displays a list of predefined graphs based on the CI that you select.
3. Select a graph or graphs from the list of Graphs. You can select more than one graph from the list by pressing the **CTRL** key while selecting the graphs.
4. Click **Draw Graphs** .

The Performance pane displays the graphs drawn.

To draw custom graphs:

1. From the **Configuration Items** tree, select the CI, for which you want to draw graphs. The **Metrics** tab on the Performance pane displays the following:
 - **Metric Classes** - Displays the list of available metric classes. This list is available only if there is more than one metric class.



- **Instances** - Displays all the available instances for a metric class. This list is available only when you select a multi-instance metric class.
- **Metrics** - Metrics pane displays the list of metrics corresponding to the selected or available metric class.

For more information regarding the vCenter metrics and different logging levels, see the *VMware documentation*.

2. Select the required metric class. The Metrics box displays the list of numeric metrics belonging to that class.
3. Drag one or more metrics from the metric list onto the Performance pane. The graph is drawn for the selected metrics.

You can drag any number of metrics and draw graphs.

To modify your selection of graph for a selected CI:

1. Select a graph type from the list of Graphs.
2. Click **Draw Graphs** . HPE Cloud Optimizer draws the selected graph.
3. Click **Clear Selections** , to cancel all the previous selections.

Saving as Favorite

To save drawn graphs as a favorite, follow these steps:

1. Draw the required graphs. For more information, see ["Drawing Graphs" on the previous page](#).
2. Click **(Save as Favorite)** in the Performance pane. The Save as Favorite dialog box opens.
3. Type the name of the group in the **Enter Favorite Name** field.

Note: If you are adding new graphs to an existing favorite list, select the required Favorite from the drop-down list.

4. Click **Save** to save the favorite list.
Click **Cancel** if you do not want to save the graphs to any lists.

Deleting a Favorite

To delete a saved favorite, follow these steps:

1. Click **(Delete Favorites)** from the Favorites pane. The Delete Favorites dialog box opens.
2. Select the favorite to be deleted from the list of favorites available in the **Select**

Favorites Name(s) box.

To select multiple favorites, press the **CTRL** key and select the required favorites.

3. Click **Delete**. The system displays a confirmation message asking to delete the selected favorites.
4. Click **Yes** to confirm.

The Favorites pane does not show the removed favorites.

Drawn Graphs Functions

Performance Pane shows the drawn graphs. Following are the functions you can perform in the graph windows:

Compare Metrics from Different CIs

1. Select a CI. The available metrics and predefined graphs for the chosen CI appears. Drag a metric from the Metrics pane. Performance Pane shows the drawn graph.
2. Select the second CI with which you want to compare the data for the first CI. Select the same metric from the metrics pane. Drag the metric to the previously drawn graph to compare the data across CIs.

Remove metric

Right-click the metric name from the legend and select **Remove**. The line representing the removed metric does not appear in the graph and the metric name does not appear in the legend.

Add metric

Drag a metric into a drawn graph window. The new metric appears in the graph and is listed in the legend.

Note: When you add or remove a metric for a default graph, the name of the graph changes as it is no longer a default graph and becomes a custom graph.

Hide metric

Click the metric name from the legend in the graph window. The metric legend is hidden and the graph of the metric does not appear in the graph window.

Show metric

Click the disabled metric from the legend. The window shows the metric legend and the graph of the metric appears on the window.

Copy metrics from one graph to another

You can drag a metric from one graph and add it to another graph. The target graph displays the newly added metric.

Pull out metric from a graph

You can drag a metric from one graph to the empty space in Performance pane. Workbench draws a new graph for the metric selected.

Move graphs

Click the top bar of the graph and drag to move or rearrange the order of the graphs drawn.

Resize graphs

Place the cursor at the bottom right corner of the drawn graph window and click and drag to resize the window.

Reports Overview

In the Reports tab, you can view the reports generated for the resources being monitored. The Reports pane displays the report for a selected CI and you can view the reports for a day, a week, or a month. For more information on viewing reports, see ["Viewing Reports" on the next page](#).

Note: Reports are available only when you install the Evaluation or Permanent License for HPE Cloud Optimizer.

Click to refresh the list of reports available for the selected CI.

The following table explains the functions available in the Reports tab.

Option	Description
Day	Shows the report for the last one day.
Week	Shows the consolidated report for the last one week.
Month	Shows the consolidated report for the last one month.
HTML	Shows the report in HTML format.
PDF	Shows the report in PDF format. You can save this report and use it later.
(Open in a	Opens the report, in the specified HTML or PDF format, in a new

Option	Description
new window)	browser window. This option is useful when you need to compare multiple reports online.

Viewing Reports

You can view the reports corresponding to a resource in the Reports pane. To view a report, follow these steps:

1. Select the required CI from the Configuration Items tree.

If you want to filter based on the type of resource and then select, click **(Filter)** and then select the required value from the list. The Configuration Items tree refreshes to display only the resources corresponding to the value set in the filter.

Note: The order of elements in the Filter By list, in the Configuration Items tree of Workbench page, is not consistent. The order changes when launching the filter list each time.

A search in the Configuration Items tree for elements starting with numbers may fail.

If a Configuration Item (CI) is renamed or a VM is moved to a different cluster, historical data for the CI is lost.

The Reports tab displays the list of available reports for the selected CI.

2. Select the required report type.

The Reports pane displays the report for the selected resource. By default, the pane displays the report generated for a day.

3. Select **Week** or **Month** to view the consolidated reports for a week or a month.
4. Select **HTML** or **PDF** to view the report in the required format. HTML is the default value.

Note: Some PDF reports in the Russian environment contain text with irregular spacing. This issue is only visible if you open the reports in Internet Explorer 8 in the Microsoft Windows Server 2008 R2 operating system.

5. Click the required type of report to view from the list of reports in the **Available Reports** pane.
6. Click **(Open in a new window)** to open the report in a new window.

Note: Some localized reports contain texts that are not localized.

vApps and folders are not available in the Treemap and Configuration Items tree, but the corresponding VMs are added to the parent in the hierarchy.

Types of Reports

HPE Cloud Optimizer provides different reports to help you analyze the performance of your virtualization environment elements. There are specific reports available for different CIs. The **Reports** tab lists the reports available for a selected CI in the element hierarchy. Some of the reports are as follows:

- **Performance Reports:** Shows the performance data.
- **Status Reports:** Shows the status summary.
- **Configuration Reports:** Shows the set up and configuration details.
- **Cross Tabulation Reports:** Shows two-dimensional data for two or more attributes and also provides a side-by-side comparison of data of your data centers.
- **Distribution Chart Reports:** Show the utilization of different resources.
- **Consolidated Reports:** Show the consolidated details for different resources.
- **Summary Table Reports:** Show a summary of the Operating Systems used in a data center.
- **Uptime Reports:** Show the uptime of the resources for the specified time period.
- **Top Instance Reports:** Show the top 10 resources in a CI.
- **Host Distribution Reports:** Show the host details for the hosts present in the cloud.
- **Usage Hours Report:** Show the usage hours for each CI.

Note: This report is available only for OpenStack.

Note: For some reports, notes are available at the end of the reports to help you interpret the reports easily.

For data to appear in the Guest Consolidated Reports, there must be data collected for at least two hours.

Using HPE Cloud Optimizer Reports for vSphere

Using HPE Cloud Optimizer, you can view different set of reports for each of the following resources in the vSphere domain:

- "Datacenter" on the next page
- "Cluster" on page 135
- "ESX/ESXi Host" on page 136

- "Resource Pool" on page 137
- "Datastore" on page 138
- "VM" on page 138

Datacenter

The following table lists the available reports and their usage for Datacenter:

Name	Description	Suggested Usage*
DataCenterResourceOverview	List of hosts, VM power status, storage, CPU, and memory allocation of the resources in the Datacenter.	<ul style="list-style-type: none"> • To view CPU and Memory allocation across the Datacenter. • To view VM activity status. • To reclaim memory allocated to idle VMs.
VMwareToolsStatus	Status of VMware Tools on each VM in a Datacenter.	To ensure compliance with VMware tools.
StorageOverviewCrossTab	Storage space utilization of the Datacenter against the available capacity.	To view storage space used and provisioned across the Datacenter and also for individual datastores.
DataCenterStatsCrossTab	Count of the VMs that have been created, destroyed, migrated, or reconfigured.	To view operational status report for the Datacenter.
DataCenterStatistics	Statistical information about the migrations and change	<ul style="list-style-type: none"> • To view Datacenter

Name	Description	Suggested Usage*
	operations about the various VMs in the Datacenter.	activity report. <ul style="list-style-type: none"> • To view the count of vMotions. • To view the type of vMotions.
DatacenterSummaryTables	Comprehensive information on server models and hardware configuration of VMs, distribution of ESX and ESXi hosts, and list of top five operating systems available in the Datacenter.	To view the physical specifications of a Datacenter.

Cluster

The following table lists the available reports and their usage for cluster:

Name	Description	Suggested Usage*
Status	Status summary of the cluster.	<ul style="list-style-type: none"> • To view CPU, Memory, and Network utilization for all the hosts in the Cluster. • To view memory overhead for individual hosts and resource pools.
HostDistribution	Host distribution of the cluster in terms of CPU and Memory utilization, and disk and network IO rate.	To view resource usage comparison.
ClusterConsolidated	Consolidated report on CPU and Memory utilization and top	<ul style="list-style-type: none"> • To view cluster activity report.

Name	Description	Suggested Usage*
	ten guests in the cluster in terms of CPU utilization.	<ul style="list-style-type: none"> To view top ten resource consumers of the cluster.
Configuration	Configuration summary of the cluster along with the effective CPU and memory utilization details.	To view effective CPU and Memory utilization.
NumberOfGuestVmotions	Count of vMotions per VM.	To view the Operational status report for the cluster.

ESX/ESXi Host

The following table lists the available reports and their usage for ESX/ESXi Host:

Name	Description	Suggested Usage*
GuestCPUReadyUtilandDemand	Utilization percentage of ready physical CPUs and the CPU demand values for each VM in the host.	To identify VMs with high CPU Ready percentage.
Status	Host details and power status of all the VMs in the host.	To reclaim the VMs that are mostly Powered Off or Powered On for longer periods.
Configuration	Configuration summary of the host as well as of all the VMs in the host.	To view complete specification of the host and individual VMs.
OverCommitStatus	Memory and CPU commitments of each VM and CPU Ready percentage for each VM.	To view CPU and Memory overcommitments for each VM.

Name	Description	Suggested Usage*
VMsWithGuestMemoryReservation	VMs with all the guest memory reserved to prevent memory over commit.	To identify VMs with guest memory reserved but not utilized.
ListOfIdleVMs	Idle VMs in the host along with their physical CPU utilization percentage and the number of CPU cores.	To reclaim the resources utilized by idle VMs to the active VMs in the Host.
CurrentCPUCyclesAllocationofVMs	Number of CPU cycles allocated versus the number of CPU cycles utilized per VM.	To view breakup of CPU allocation across the host.
HostPerformance	Memory and CPU utilization and disk and network I/O details.	To view host performance in terms of CPU, Memory, disk, and network IO.
VMPhysCPUUtilwrthHost	Minimum, maximum, and average CPU utilization of all the VMs in the host.	To view physical CPU utilization across the host.
GuestMemoryActiveUsage	Active memory for each VM in the host.	<ul style="list-style-type: none"> To view physical memory utilized by each VM. To view VMs consuming the highest memory.

Resource Pool

The following table lists the available reports and their usage for Resource Pool:

Name	Description	Suggested Usage*
View	Details of the resource	To view memory and CPU entitlement for a

Name	Description	Suggested Usage*
Status	pool along with the power status of each VM in the resource pool.	Resource Pool.

Datastore

The following table lists the available reports and their usage for Datastore:

Name	Description	Suggested Usage*
StorageIOandUsage	Overview of the storage utilization in a datastore.	<ul style="list-style-type: none"> • To view the capacity usage report. • To view the Datastore utilization and overcommitment. • To view the guests on the Datastore.
StorageUsedbyIdleVMs	Storage utilization of idle VMs on a datastore.	To reclaim storage space from Datastores by deleting idle VMs.
DataStorePerformance	Trend of space used versus the space provisioned on a datastore ,the average total latency and IO operations per VM.	<ul style="list-style-type: none"> • To view performance of the Datastore. • To view over commit status of the Datastore.

VM

The following table lists the available reports and their usage for VM:

Name	Description	Suggested Usage*
Status	Current CPU and Memory utilization (percent) and Disk and Network Read Write (kilobyte/second) details.	To view status and performance of the VM.
Configuration	Current configuration summary of the VM.	To view CPU and memory reservation for the VM.
GuestConsolidated	Active and consumed memory (percent), CPU and Memory usage (percent), and disk and network IO rate (kilobyte/second) over the selected interval of day, month, or week.	To view memory, CPU, and network usage information for the VM.

*List of common scenarios in which this report can be useful

Using HPE Cloud Optimizer Reports for Hyper-V

Using HPE Cloud Optimizer, you can view different set of reports for each of the following resource in the Hyper-V domain:

- "Host Cluster" below
- "Host" on the next page
- "Datastore" on the next page
- "VM" on page 141

Host Cluster

The following table lists the available reports and their usage for Host Cluster:

Name	Description	Suggested Usage*
Status	Status summary of the cluster along with CPU, Memory, and network utilization of the hosts in the cluster.	To view CPU, Memory, and Network utilization for all hosts in the Host Cluster.
Configuration	Configuration summary of the cluster along with Physical CPU, CPU core	To view physical CPU and CPU core count in the Host

Name	Description	Suggested Usage*
	count, and memory details.	Cluster.

Host

The following table lists the available reports and their usage for a Host:

Name	Description	Suggested Usage*
Status	Host details and power status of all the VMs in the host.	To view VM activity status.
Configuration	Configuration summary of the host as well as of all the VMs in the host.	To view VM configurations.
TopInstance	List of top ten hosts in terms of CPU and Memory utilization.	To view top ten consumers of resources.

Datastore

The following table lists the available reports and their usage for a Datastore:

Name	Description	Suggested Usage*
StorageIOandUsage	Overview of the storage utilization in a datastore.	<ul style="list-style-type: none"> To view the capacity usage report. To view the Datastore utilization and overcommitment. To view the guests on the Datastore.
DataStorePerformance	Trend of space used versus the space provisioned on a datastore and the average total latency and IO operations per VM.	<ul style="list-style-type: none"> To view performance of the Datastore. To view over

Name	Description	Suggested Usage*
		commit status of the Datastore.

VM

The following table lists the available reports and their usage for a VM:

Name	Description	Suggested Usage*
Status	Current CPU and Memory utilization (%) and Disk and Network Read Write (KB/sec) details.	To view status and performance of the VM.
Configuration	Current configuration summary of the VM.	To view CPU and memory reservation for the VM.

*List of common scenarios in which this report can be useful

Using HPE Cloud Optimizer Reports for OpenStack

Using HPE Cloud Optimizer you can view different set of reports for each of the following resource in the OpenStack domain:

- "Cloud" below
- "Tenant" on the next page
- "Host" on the next page
- "VM" on page 143

Cloud

The following table lists the available reports and their usage for a Cloud:

Name	Description	Suggested Usage*
Summary	Cloud summary in terms of the number of VMs, Hypervisors, and tenants, top ten and bottom ten VMs by CPU utilization.	To view top ten and bottom ten VMs in terms of CPU utilization.

Name	Description	Suggested Usage*
HostDistribution	Host distribution in terms of memory and CPU utilization, number of CPU cores, and the number of VMs	To view CPU and memory utilization across all the VMs in the host.

Tenant

The following table lists the available reports and their usage for a Tenant:

Name	Description	Suggested Usage*
Summary	CPU and memory utilization, top ten and bottom ten VMs by CPU utilization.	To view top ten and bottom ten VMs in terms of CPU utilization.
Status	Tenant details and the status of virtual machines.	To view VM activity status.
Configuration	Configuration summary of the tenant as well as of all the VMs in the tenant.	To view configuration of VMs across the tenant.
Performance	Performance summary of the tenant in terms of CPU, memory, and disk space utilization.	To analyze performance of the tenant.
UsageHours	Usage summary in case of CPU, Memory, and disk space.	To view usage of CPU, Memory, and disk space.

Host

The following table lists the available reports and their usage for a Host:

Name	Description	Suggested Usage*
Summary	CPU and memory utilization, top ten and bottom ten VMs by CPU utilization.	To view top ten and bottom ten VMs in terms of CPU utilization.
Status	Host details and the status of virtual machines.	To view VM activity status.
Configuration	Configuration summary of the host as well as of all the VMs in the host.	To view configuration of VMs across the host.

Name	Description	Suggested Usage*
Performance	Performance summary of the host in terms of CPU, memory, and disk space utilization.	To analyze performance of the host.
UsageHours	Usage summary in case of CPU, Memory, and disk space.	To view usage of CPU, Memory, and disk space.

VM

The following table lists the available reports and their usage for a VM:

Name	Description	Suggested Usage*
Status	Current CPU and Memory utilization (%) and Disk and Network Read Write (KB/sec) details.	To view memory, CPU, and network usage information for the VM.

*List of common scenarios in which this report can be useful

Using HPE Cloud Optimizer Reports for KVM/Xen

Using HPE Cloud Optimizer you can view different set of reports for each of the following resource in the KVM/Xen domain:

- "Host" below
- "VM" on the next page

Host

The following table lists the available reports and their usage for a Host:

Name	Description	Suggested Usage*
Status	Status summary of the host along with CPU, and Memory utilization. You can also see the VM power status.	To view status and performance of the VM.
Configuration	Configuration summary of the host and the VMs.	To view complete specification of the host and individual VMs.

VM

The following table lists the available reports and their usage for a VM:

Name	Description	Suggested Usage*
Status	Status summary of the VM along with CPU utilization, disk and network read write details.	To view memory, CPU, and network usage information for the VM.

*List of common scenarios in which this report can be useful

Using HPE Cloud Optimizer reports for HPE OneView

When HPE Cloud Optimizer is integrated with HPE OneView, you can view reports for the following components:

- Enclosures
- ServerHardwares or Blade Servers
- VMware Cluster

Enclosures

The following table lists the available reports and their usage:

Name	Description	Suggested Usage*
Overview	General enclosure information, Hardware information, Power and Temperature utilization, number of vacant bays, distribution of virtualization types in the Enclosure	<ul style="list-style-type: none">• To get an overview of different virtualization types on the blades in an enclosure.• To view the inventory of clusters and hosts across enclosure bays.
Virtual Machine Inventory	List of VMs running on the enclosure	<ul style="list-style-type: none">• To view the inventory of VMs running on the enclosure.• To view a mapping of VMs with the Blade

Name	Description	Suggested Usage*
		<p>servers on which they are configured.</p> <ul style="list-style-type: none"> To help identify the VMs that would be affected if the specific enclosure is not available.
Utilization	Utilization summary of the enclosure in terms of CPU, Memory, and Power.	To view CPU, Memory, Power, and Bay utilization of the enclosure.
Bay Topology	Layout of device bays in the enclosure, bays with hypervisors and their specific types.	<ul style="list-style-type: none"> To identify vacant and occupied bays. To identify if the server blade is half-height or full-height. A full-height server blade occupies two device bays but is associated with the top bay only. To assess the impact on a hypervisor if the respective bay is removed.
Host Power Utilization	Performance Per Watt summary for all the hosts on the Enclosure.	To view the Power and CPU Utilization for each Host available on the Enclosure

ServerHardwares or Blade Servers

Name	Description	Suggested Usage*
Overview	Information about the hypervisor environment, configuration summary of the host and VMs configured on the Blade server	<ul style="list-style-type: none"> To view the configuration details of the VMs on the Blade server. To view CPU, Memory, and Power utilization of the hypervisors.
Utilization	Power utilization, Memory and CPU utilization, and disk and network I/O details.	To view Blade server performance in terms of Power, CPU, Memory, disk, and network IO.

VMware Cluster

Name	Description	Suggested Usage*
Risk Analysis	This report appears in the vSphere domain. It gives general information about cluster, and layout of the cluster.	<ul style="list-style-type: none">• To view the layout of the cluster across enclosures and standalone hosts.• To assess if a cluster can sustain itself when one of the enclosures on which the cluster is configured is shut down for maintenance. <p>For example, a cluster spans across an enclosure and a standalone host. If the enclosure has to be shut down for maintenance, this report tells you if the cluster VMs on that enclosure can be successfully migrated to the standalone hosts without affecting the cluster.</p> <p>The cluster is considered to be at:</p> <ul style="list-style-type: none">• High risk if the Days to Capacity (DTC)¹ is less than or equal to two.• Medium risk if Days to Capacity (DTC) is between two and five.• Low risk if the Days to Capacity (DTC) is greater than five.

*List of common scenarios in which this report can be useful

Filter

The Performance Filter allows you to filter domain specific entities based on metric values. You can add multiple filters for the entities. The filters are based on the rules applied to the metrics.

To filter entities based on metric values, follow these steps:

1. In the **Domain** drop-down, select the domain.
2. In the **Entity Type** drop-down, select the type of entity to be filtered.
3. In the **Filter** section, select a metric name.
4. In the **rule** drop-down, select a relational operator.

¹the number of days in which the entity reaches the maximum capacity

5. Enter a value for the metric based on the relational operator selected.
6. Click **Add Filter** to add multiple filters for the entity type.
7. Click **Submit**. A list of entities appear that match the entered filter criteria.

Chapter 7: Alerts

In the Virtualization environment, as an administrator, you may observe the following conditions:

- Sudden capacity violations and performance problems
- Rapid growth or drop in trends of resource usage
- Reduction in capacity of your IT infrastructure in the perspective of business services, applications, and clusters

These situations can result into poor performance and must be avoided. Avoid this situation by receiving the smart detailed alerts much before the conditions occur in the virtualized environment.

Alerts¹ enable you to quickly identify and troubleshoot the problems in the virtualized environment. Alert analysis starts with the data collection.

Alerts allow you to perform the following tasks:

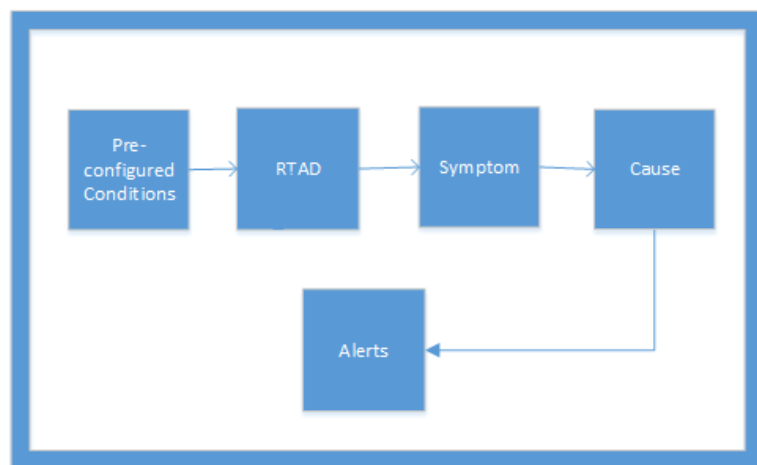
- Detect the symptom in the virtualized environment
- Find the cause of symptom
- Evaluate the root cause based on the pre-defined conditions

The content- rich alerts help you fix the problem before it occurs.

Whenever there is a sudden change detected in the environment, the **Real Time Alert Detection (RTAD)** works on the pre-configured conditions available to find the symptom. A symptom is a message generated to notify the occurrence of abnormal behavior in the environment. Analysis is done based on the default settings available to find the cause of the symptom. Further drill down is done to find the actual root cause that triggered the symptom. If the cause persists for a certain time period in the environment, an alert is generated. An alert is generated based on the type of **sensitivity** deployed. The type of sensitivity deployed defines the time to react to a given situation and issue an alert. Only when the cause and root cause is detected, then the smart alert with detailed description about the problem is available on the HPE Cloud Optimizer console.

Working of Smart Alert Framework (SAF)

¹Messages that provide structured information about the problem.



You can assess the health of the environment by checking the HPE Cloud Optimizer console to identify areas that require your attention.

Example:

In a virtualized environment, you can use dynamic methods to maintain an optimal balance between the demand and utilization of resources, thus avoiding over-utilization of resources. Also, you would need proper notification if there is a problem in the environment.

Scenario: Host shows low CPU availability

In a virtualized environment, CPU utilization can be high for a particular period of time but this condition may get resolved and hence does not qualify to become a symptom. There can be various reasons for this condition.

If the issue persists for some time then a symptom is generated. To find the cause, check if the host is saturated because VMs are misconfigured or the workload is high. You can drill down further and conclude the root cause of the issue.

When the root cause is detected for the symptom, a smart alert is generated. You can view the alert from the Alert Visualizer on the HPE Cloud Optimizer console.

Baseline Alerts

The alerting component also generates the **Baseline Alerts**. The baseline alerts provide information about the significant deviation in the real time situation. These baseline alerts are generated only after data is collected for 10 samples. The baseline alerts work on the collected sample data to analyze the trends and if there is any deviation that is observed then only baseline alerts are generated. If the baseline value is below the DoNotBotherBelow parameter, the baseline alert will not be generated. If the baseline value shows a sharp increase or is above the DoNotBotherBelow parameter value, the baseline alert will be generated.

These baseline alerts can help to preempt the performance of virtualized environment and also evaluate false alerts.

Note: Baseline alerts are generated only for Disk and Memory.

Example:

In a virtualized environment, consider memory utilization of the host to be around 30% - 40% for a particular period of time. The baseline is calculated with the last 10 collected samples. Baseline alerts are generated if the current memory utilization is higher than the calculated baseline or if there is a sharp increase in the memory utilization.


Using Alert Messages

Navigation

To access the feature, from the HPE Cloud Optimizer home page, select **Performance > Alerts**. The Alert Messages page opens.

You can also access the page from Dashboard. To get the list of alerts, in the Host(s), VM(s), or Datastore(s) boxes, click Critical or Major. The respective Alert Messages page appears.

The following table lists the tasks you can perform from the Alert Visualizer:

Tasks	Description
View Capacity Data	Select an alert in the alert table. From the Tools menu, select  to go to the entity Summary page to view the capacity and usage data of the particular entity. For more information, see "Using Capacity" on page 170 .
View Cause of the Alert	Select the alert description to view the root cause of the problem. You can view the hosts affected and view the reports.
Performance Graphs Overview	View or draw graphs for the performance data of the particular entity.
Filter Alerts based on the requirement	View alert(s) based on the specific entity, description, Domain, and Entity Type.
View Alerts	View alert(s) generated in the specific time period.

Tasks	Description
for a specific time	

Alert View

The following table provides information on the user interface elements on the page:

Element	Description
Entity	Name of the Entity. You can sort the entities in the ascending or descending order by clicking the Entity field.
Severity	Lists the severity of the message. The values for Severity are Critical, Major, Minor, Warning, and All.
Time	Displays the time and date when the alert was generated.
Description	Displays the root cause of the alert.
Domain	Displays the domain of the entity.
Entity Type	Represents the type of entity.
(Sort the columns)	Use this option to sort the data in the columns in the ascending or descending order.
Search	Type a value in this field to search for specific events based on Entity, Severity, Description, Domain, and Entity Type.
Refresh	Refreshes the list of entities and alerts displayed on the page.
Filter	Displays the alerts specific to the Entity, Severity, Time, Description, Domain, and Entity Type. For Example: Apply Filter on Domain and type VCENTER to view all the alerts belonging to the VCENTER domain.

Element	Description
Select Time Range	Displays the alerts generated during that time range. The options available are as follows: <ul style="list-style-type: none">• Last Hour• Last 6 Hours• Last 12 Hours• Last One Day

Conditions Used to Generate the Alert Messages in VMware

You can assess the health of the environment by checking the HPE Cloud Optimizer console to identify areas that require your attention. Alerting framework monitors the VMware virtualized environment for the following conditions

For CPU:

Definition	Condition	Possible causes
Host CPU is saturated	CpuReadyUtil > Threshold, and CpuDemand > Capacity - HeadRoom	Multiple VMs with High CPU workload demand running on this Host.
Host CPU is contended	CpuReadyUtil > Threshold, and there is no significant CpuDemand from a large number of Virtual Machines running on the Host.	Multiple VMs running on this Host.
Virtual Machine CPU is Saturated	CpuReadyUtil > Threshold, and CpuUtil > 100 – HeadRoom	VM needs more CPU than currently allocated.
Virtual Machine CPU	CpuReadyUtil > Threshold, VM CPUUtil < HeadRoom %	VMs are configured with more number of vCPU(s) and there is a possibility that other VMs on the host are causing CPU

Definition	Condition	Possible causes
parameters are not configured properly		contention.
Cluster CPU is saturated	All Hosts in the cluster are saturated	<ul style="list-style-type: none"> VMs are actively running on all the hosts Multiple VMs(where CPU utilization is high) are running on the Hosts at the same time
Abnormal high number of vMotions.	vMotions > VMotionThreshold (Baselined)	<ul style="list-style-type: none"> The DRS Cluster may have large number of Virtual machines that are contending for CPU or Memory resource and all the hosts in the Cluster is in resource saturated condition. DRS mode is set to Aggressive and/or Auto. A host in a DRS cluster is put in maintenance mode causing all the high available VMs to migrate to other Host in the same cluster. Hosts in the cluster is not running balanced CPU and Memory workload due to affinity rules applied at each Virtual Machine level which is causing other VMs to vMotion.

For Memory:

Definition	Condition	Possible causes
Host memory is saturated.	MemoryBallooned > 0, MemActive > (Capacity – HeadRoom %) and MemActive has Increasing trend	<ul style="list-style-type: none"> Multiple VMs with high memory workload are running on the Host. Guest tool status on virtual machines is not in one of the recommended states. The not recommended states are: Not Installed, Not Running, Out Of Data, or Unknown. Large number of VM(s) configured with Memory reservations. Too many memory reservations may result in memory

Definition	Condition	Possible causes
		saturation.
Abnormal high memory usage on Host detected.	MemUtil > Threshold	<ul style="list-style-type: none"> Multiple VMs with High memory demand are running on the Host. Abnormal increase in Memory usage may be due to additions of Virtual machines due to manual or DRS initiated vMotions.
VM memory configuration is not set for optimal usage of Host resource.	Memory Reservation is set And MemActive < 50% of Memory Reservation	NA
VM memory allocation is in surplus	MemBalloonUsed > 0 And MemActive < MemHeadRoom(20)% of MemConfigured on VM	NA
VM memory is saturated	Memory Limit is set Memory Balloon > 0 or Swap Swapped > 0 MemActive > 95% of Memory Limit	Memory usage limit is set on the VM is causing huge performance impact on the behavior.
Cluster memory is saturated	All Hosts in the cluster are saturated by Memory resource.	<ul style="list-style-type: none"> Multiple VMs are running memory intensive workload on the cluster. Each Host in the cluster has reserved memory capacity because of VM memory reservations, and hence memory reclamation techniques (ballooning) cannot be applied to reclaim un-used memory. Cluster is out of Memory capacity to support the memory demand of all Virtual machines running at a given point in time. <p>Memory saturation at the cluster may result in multiple vMotions if DRS is enabled.</p>

Definition	Condition	Possible causes
Abnormal high number of vMotions.	vMotions > VMotionsThreshold (Baselined)	<ul style="list-style-type: none"> The DRS Cluster may have large number of Virtual machines that are contending for CPU or Memory resource and all the hosts in the Cluster is in resource saturated condition. DRS mode is set to Aggressive and/or Auto. A host in a DRS cluster is put in maintenance mode causing all the high available VMs to migrate to other Host in the same cluster. Hosts in the cluster is not running balanced CPU and Memory workload due to affinity rules applied at each Virtual machine level which is causing other VMs to vMotion.

For Datastore:

Definition	Condition	Possible Causes
Low free space on Datastore	Free Space < HeadRoom	<ul style="list-style-type: none"> Large number of VMs with thick provision disks on this Datastore. Too many aged snapshots. Large number of VMs deleted from vCenter inventory but VMDK (s) and Snapshot files retained in the datastore for those VMs. Large number of idle VMs on the datastore.
Virtual machine is observing abnormal IO activity	Disk I/O Rate is abnormal and Disk I/O Rate > 500 Kilobytes per second (kbps)	<ul style="list-style-type: none"> The workload on the VM is sending too many IO requests to the Disks. For example, a software backup or antivirus on-demand scan is running during this time. Applications that can utilize system memory to cache data and avoid disk access can also result

Definition	Condition	Possible Causes
		<p>in high IO(s) when memory usage on the VM is at peak.</p> <ul style="list-style-type: none"> • VM is swapping to their swap disks which is causing High I/O on the disks. • Large number of snapshot disks for a VM could also result in high disk IO activity.
High Disk IO Latency has been observed	Disk Kernel Latency > KernelLatencyThreshold	VMs on the host are trying to send more throughput to the storage system than the configuration support.

Conditions Used to Generate the Alert Messages in Hyper-V

You can assess the health of the environment by checking the HPE Cloud Optimizer console to identify areas that require your attention. Alerting framework monitors the Hyper-V virtualized environment for the following conditions:

For CPU:

Definition	Condition	Possible causes
Host CPU is saturated	HostCPUUtilization > HostCPUUtilThreshold	Multiple Virtual Machines actively running on this Host.
Abnormal high CPU usage is detected		<ul style="list-style-type: none"> • Multiple Virtual Machines with High CPU demand is running on this Host. • Addition of Virtual machines due to manual or live migration.
Virtual Machine CPU is Saturated	CPUCycleTotalUsed > 95% of CPU Limit	The Virtual Machine CPU Usage has reached up to 95% or more of the set CPU limit.
Virtual Machine CPU is saturated	CPU Utilization > VMCPUUtilizationThreshold	Virtual machine is running with a CPU intensive workload.

For Memory (Hosts):

Definition	Condition	Possible causes
High Memory usage is detected	Memory Utilization > HostMemoryUtilizationThreshold	<ul style="list-style-type: none"> Multiple Virtual Machines with High Memory demand are running on this Host. High Memory usage may also be caused due to addition of Virtual machines through manual or live migration.
Abnormal high Memory usage is detected		<ul style="list-style-type: none"> Multiple VMs with High memory demand are running on the Host. Abnormal increase in Memory usage may be due to addition of Virtual machines through manual or DRS initiated vMotions.
Host Memory is saturated	MemoryPressure > 100, HostMemDemand[MB] > (Capacity [MB] - Memory HeadRoom[%] - Total Configured Static Memory	<ul style="list-style-type: none"> Multiple Virtual Machines with high Memory workload are running on this Host. Multiple Virtual machine(s) configured with high Minimum Memory.

For Memory (VMs):

Definition	Condition	Possible causes
Virtual Machine Memory is saturated	Memory Demand [MB]> Memory Consumed[MB]	<ul style="list-style-type: none"> Virtual machine is running with a memory intensive workload. Multiple Virtual Machines with high Memory weight are running on this Host.

For Datastore:

Definition	Condition	Possible Causes
Low free space on Datastore	Free Space < HeadRoom	<ul style="list-style-type: none"> Large number of VMs with thick provision disks are available on this Datastore. Multiple aged snapshots. Large number of VMs deleted from the inventory

Definition	Condition	Possible Causes
		<p>but Virtual Machine Disk Files (VMDKs) and Snapshot files for those VMs are retained in the datastore.</p> <ul style="list-style-type: none"> • Large number of idle VMs on the datastore.

Conditions Used to Generate the Alert Messages in KVM

You can assess the health of the environment by checking the HPE Cloud Optimizer console to identify areas that require your attention. Alerting framework monitors the KVM virtualized environment for the following conditions:

For CPU:

Definition	Condition	Possible causes
Virtual Machine CPU is saturated	CPU Utilization > VMCPUtilizationThreshold	Virtual machine is running with a CPU intensive workload.

For Memory (VMs):

Definition	Condition	Possible causes
Virtual Machine Memory is saturated	Active Memory > Memory Threshold	<ul style="list-style-type: none"> • Virtual machine is running with a memory intensive workload. • Multiple Virtual Machines with high Memory weight are running on this Host.
Abnormal high Memory pressure is detected	Abnormal high memory pressure is detected	<ul style="list-style-type: none"> • Multiple Virtual Machines with High Memory demand is running on this Host. • Additions of Virtual machines due to manual or through live migration.
High Memory usage is detected	Memory Utilization > HostMemoryUtilizationThreshold	<ul style="list-style-type: none"> • Multiple Virtual machines with High Memory demand

Definition	Condition	Possible causes
		is running on this Host. <ul style="list-style-type: none">• Additions of Virtual machines due to manual or through live migration.

For Datastore:

Definition	Condition	Possible Causes
Low free space on Datastore	Free Space < HeadRoom	<ul style="list-style-type: none">• Large number of VMs with thick provision disks are available on this Datastore.• Large number of idle VMs on the datastore.

Monitoring Alerts on the HPOM Console

HPE Cloud Optimizer integration with HPOM helps you to monitor the abnormalities available in the virtual infrastructure from an HPE Operations Manager (HPOM) console. You may monitor the performance, capacity, utilization and resource consumption of the host machines and virtual machines from the HPOM console. HPE Cloud Optimizer provides the capability of analyzing and generating the smart alerts for abnormalities available in the environment. HPE Cloud Optimizer forwards the alerts to HPOM console. You can view the smart alerts on the HPE Cloud Optimizer console and HPOM console.

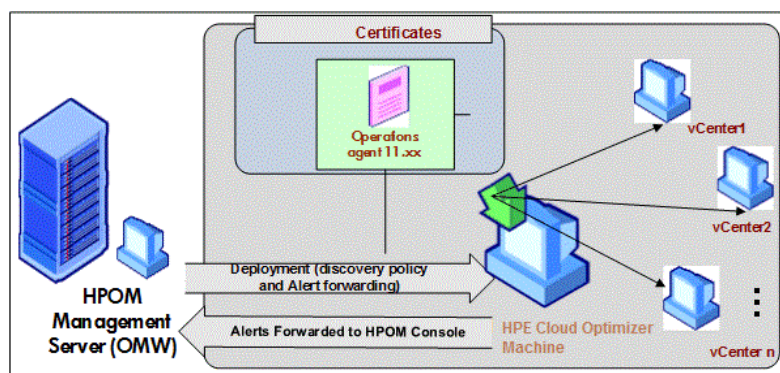
Additional benefits for integrating with HPOM:

- Easy discovery and monitoring of VMware data center topology
- Quick view of the smart alerts on HPOM console

Integration with HPOM

To integrate HPE Cloud Optimizer with HPOM, you must install HPE Operations agent. The HPE Operations agent installed on a HPE Cloud Optimizer machine can monitor the environment by collecting the performance data from individual ESXi hosts. This type of monitoring requires you to install and configure the HPE Operations agent on the HPE Cloud Optimizer machine. Deploy the HPE Cloud Optimizer Integration policies to forward the HPE Cloud Optimizer alerts to HPOM server.

The following illustration shows HPE Cloud Optimizer integration with HPOM:



To integrate HPE Cloud Optimizer integration with HPOM, follow the [steps](#). To view the details about alerts, see [Managing Alerts](#).

Configuring Sensitivity Levels in HPOM Machine

To configure sensitivity levels in HPOM machine, perform the following steps:

1. Log on to the HPOM machine.
2. In the left pane, from the **Policy Management Group**, select **vPV-Integration**.
3. Open the **vPV-AlertSensitivity** policy.
4. In the **vPV-AlertSensitivity** configuration file, set the sensitivity level as Low, Medium, High, or Custom.
5. If the sensitivity level is set as **Custom**, open the **vPV-CustomAlertSensitivityDefinition** policy.
 - a. Edit the sensitivity parameters as per your requirements.
 - b. Deploy the **vPV-CustomAlertSensitivityDefinition** policy to HPE Cloud Optimizer machine.
6. Deploy the **vPV-AlertSensitivity** policy to HPE Cloud Optimizer machine.

Sensitivity

In a virtualized environment, you must get a notification when an issue occurs in the environment. To get these notifications, you may have to evaluate the available conditions that may exist in the environment and then set the threshold values. As an administrator, with this approach, you may need additional time and resources to complete the task.

Sensitivity is based on the set of rules where the threshold range and values are defined for all the available resources. The generic set of rules are defined for CPU, memory, and data sources that contain the pre-defined threshold values. An alert is generated based on the type of sensitivity deployed. The type of sensitivity deployed defines the time to react to a given situation and issue an alert.

The sensitivity is categorized as follows:

- Low - This type sends the alert when the alert type situation persists for more than 20 minutes.
- Medium - This type sends the alert when the alert type situation persists for more than 10 minutes.
- High - This type sends the alert immediately as there is no wait time.
- Custom - You can customize the pre-defined threshold values to generate the alerts.
- Disable - This option is available to stop the alerts.

The default sensitivity deployed is medium. You can configure the sensitivity based on the requirement.

Configuring Sensitivity

- **To change the sensitivity as Low, High, Medium, or Disable**

1. Log on as an administrator.
2. Run the script: `/opt/OV/sbin/PV/ChangeAlertSensitivity.py`
`[sensitivity]`

In this instance,

the parameter `sensitivity` can have the values as High, Low, Medium, or Disable.

You can enter the type of sensitivity as per your requirement. Example: If you change the sensitivity to High, then when the symptom occurs, the alert will be generated immediately.

Note: Select `Disable` option to disable the alerting feature.

- **To change the sensitivity as Custom**

1. Log on as an administrator.
2. Go to the following location:
`/var/opt/OV/conf/PV/Monitoring/SensitiveFiles`
3. Open the `vPV_Alerts_Custom_Configuration.ini` file.
4. Edit the sensitivity parameters as per your requirements.

Note: Additionally, you can change the threshold value for datastore space utilization. The threshold value for datastore space utilization = `100 - RIGHTSIZE_DISK_HEADROOM`. By default, the parameter `RIGHTSIZE_DISK_HEADROOM` is set to 10. So, the default threshold value is 90.

To change the threshold value to 75,

- a. Edit the `vPV-CustomAlertSensitivityDefinition` policy.

b. Update `RIGHTSIZE_DISK_HEADROOM` as follows:

```
[pvcd.P0]  
RIGHTSIZE_DISK_HEADROOM=25
```

5. Run the script: `/opt/OV/sbin/PV/ChangeAlertSensitivity.py custom`.
Alerts are now be sent based on the customized threshold values.

Suppressing Alerts

While administering virtualized environments, you often come across scenarios where there is a need to suppress alerts from certain virtualized entities. The ability to suppress alerts from specific entities can help reduce alerting noise during a planned maintenance activity of the virtualized environment.

In HPE Cloud Optimizer, you can suppress alerts for VMs, Hosts, Clusters, and Datastores based on user-defined rules and conditions. If HPE Cloud Optimizer is integrated with HPE Operations Manager (HPOM) or HPE Operations Manager *i* (HPE OMi), you can suppress the alerts directly from the HPE OM and HPE OMi.

Suppressing alerts in HPE Cloud Optimizer

To suppress the alerts in HPE Cloud Optimizer, perform the following steps:

1. On the HPE Cloud Optimizer machine, log on as root.
2. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

In this instance, `install_dir` is the directory where HPE Cloud Optimizer is installed. The default location is `/opt/OV`.

A text file opens.

3. In the file, type the following to create a namespace:

```
[pvcd.alert.suppress.<Entity_Type>.rules]
```

In this instance, `<Entity_Type>` is `vm`, `host`, `datastore`, or `cluster`. The `Entity_Type` is the entity for which you want to suppress the alerts. If you want to suppress alerts for multiple entity types, you need to create namespace for each entity type.

For example, if you want to suppress alerts for VM and Host, create the following two namespaces:

```
[pvcd.alert.suppress.vm.rules]
```

```
[pvcd.alert.suppress.host.rules]
```

Note: All the values in the text file are case sensitive.

4. Under the [pvcd.alert.suppress.<Entity_Type>.rules] namespace, type the following to define the number of Rules that would be used to suppress the alerts for the specific entity:

```
rules=<RuleName1>,<RuleName2>,...<RuleName-n>
```

In this instance,

RuleName1, *RuleName2*, *RuleName-n* are the names assigned to the rules. The names assigned here must be used as it is in the subsequent rule declaration namespace.

n is the total number of rules that should be defined to suppress the alerts for the specific entity.

5. Create namespaces equivalent to the number of Rules defined in the [pvcd.alert.suppress.<Entity_Type>.rules] namespace. All the Rules defined in Step 4 should be declared using the following namespaces:

```
[pvcd.alert.suppress.<Entity_Type>.<RuleName1>]
```

```
[pvcd.alert.suppress.<Entity_Type>.<RuleName2>]
```

...

```
[pvcd.alert.suppress.<Entity_Type>.<RuleName-n>]
```

In this instance,

RuleName1, *RuleName2*, *RuleName-n* are the rule names defined in the [pvcd.alert.suppress.<Entity_Type>.rules] namespace.

6. Under each [pvcd.alert.suppress.<Entity_Type>.<RuleName-n>] namespace created above, define clauses that should be satisfied for a rule to be implemented for alert suppression.

Note: A rule can be applied only if all the clauses in the rule are satisfied.

```
[pvcd.alert.suppress.<Entity_Type>.<RuleName>]
```

```
Clause1=<Metric_Name>:<Operator>:<Value>
```

```
Clause2=<Metric_Name>:<Operator>:<Value>
```

..

```
Clause-n=<Metric_Name>:<Operator>:<Value>
```

In this instance,

- *Clause1*, *Clause2*, .., *Clause-n* are the names assigned to the clauses.
- *n* is the total number clauses that are defined.
- *Metric_Name* is the name of the metric to be compared in the rule.

- *Operator* is a symbol denoting an operation. Operator is used to compare values.

Following is a list of supported operators:

Operator	Description
EQ	Equals
NE	Not Equal To
GT	Greater Than
GTE	Greater Than or Equal To
LT	Less Than
LTE	Less Than or Equal To
LIKE	Search for a specified pattern in the value

- *Value* is the metric value that is compared using the specified operator.

7. Save and close the file.

Suppressing HPE Cloud Optimizer alerts in HPE OMi

To suppress HPE Cloud Optimizer alerts in HPE OMi, perform the following steps:

1. Open the **Management Templates and Aspects** manager:
Admin > Operations Management > Monitoring > Management Templates and Aspects
2. In the **Configuration Folders** pane, click **Configuration Folders > Infrastructure Management > vPV Management Pack**.
3. In the **Management Templates and Aspects** pane, select the **vPV Alerts** Aspect.
4. Right-click and select **Edit Item**. The **vPV Alerts: Edit Aspect** window appears.
5. Click the **Policy Templates** tab. A list of existing Policy Templates appears.
6. Double-click the **vPV-SuppressAlerts** Node Info Template. The Node Info property window appears.
7. Click the **Policy Data** tab to get the Node Info Policy view.
8. Edit the node info policy as specified in steps 2-6 of the [Suppressing alerts in HPE Cloud Optimizer](#) section.
9. Click **OK** to save and close the file.
10. Deploy the **vPV-SuppressAlerts** policy to HPE Cloud Optimizer machine.

Suppressing HPE Cloud Optimizer alerts in HPOM

To suppress HPE Cloud Optimizer alerts in HPOM, perform the following steps:

1. Log on to the HPOM machine.
2. In the left pane, navigate to **Policy management > Policy group**.
3. In the **Policy group** tree, select **vPV-Integration**. The various policies appear in the right-pane.
4. In the right-pane, select the **vPV-SuppressAlerts** policy.
5. Edit the **vPV-SuppressAlerts** node info policy as specified in steps 2-6 of the [Suppressing alerts in HPE Cloud Optimizer](#) section.
6. Save and close the file.
7. Deploy the **vPV-SuppressAlerts** policy to HPE Cloud Optimizer machine.

Example

Adam is a virtualization Subject Matter Expert in an organization. He has planned for a maintenance activity in his virtualized environment. During that time, Adam does not want to receive alerts from the test VMs and Hosts in his environment.

Adam wants to suppress alerts from VMs and Hosts that have 'test' as part of the system name. To suppress alerts from all the test VMs and Hosts, perform the following steps:

1. Log on to the HPE Cloud Optimizer machine as root.
2. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

In this instance, `install_dir` is the directory where HPE Cloud Optimizer is installed. The default location is `/opt/OV`.

A text file opens.

3. In the file, type the following:

```
[pvcd.alert.suppress.vm.rules]
rules=Rule1
[pvcd.alert.suppress.host.rules]
rules=Rule1
[pvcd.alert.suppress.vm.Rule1]
Clause1=SystemName:LIKE:%test%
[pvcd.alert.suppress.host.Rule1]
Clause1=SystemHostName:LIKE:%test%
```

4. Save and close the file.

Monitoring the State of Virtual Machines

While monitoring a virtualized environment, it is important to receive alerts when there is an unexpected change in the state of the Virtual Machines. With HPE Cloud Optimizer, you can receive alerts when there is a change in the state of the Virtual Machine. You can configure the VM states for which you would like to receive alerts. You can also configure and stop the alerts from being generated for some VM states, for example, when the VMs are a part of a planned maintenance activity.

Note: VM state change alerts are generated only for VMs in the vCenter, KVM, and Hyper-V domain.

To configure the VM states for which you want receive alerts, perform the following steps on the HPE Cloud Optimizer machine:

1. Log on to the machine as root.
2. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

In this instance, `install_dir` is the directory where HPE Cloud Optimizer is installed. The default location is `/opt/OV`.

A text file opens.

3. In the file, locate the `[pvcd.alert.<Virt_Type>.StateMonitor]` namespace. In this instance, `Virt_Type` is virtualization domain for which you want to generate the VM state change alerts. `Virt_Type` can be VMWARE, KVM, and HYPERV. For example, `[pvcd.alert.VMWARE.StateMonitor]`.

Note: All the values in the text file are case sensitive.

The `[pvcd.alert.<Virt_Type>.StateMonitor]` namespace has the following parameters:

```
[pvcd.alert.<Virt_Type>.StateMonitor]
AlertOnDownStates=off
ShortTermPeak=0
AbnormalStates=stuck,crash,hung,unknown,invalid,shutdown,other
DownStates=paused,off,boot,suspended,shutdown,disabled,down
NormalStates=on,up,enabled,running,active
```

The following table contains information on how to configure the various parameters:

Parameter	Description
AlertOnDownStates	<p>By default, the alerts are not generated when the VM state changes to any of the VM down states configured in the <code>DownStates</code> parameter list. Alerts are only generated for the VM states configured in the <code>AbnormalStates</code> parameter list.</p> <p>To start receiving alerts when the VM state changes to any of the VM down states, set the <code>AlertOnDownStates</code> parameter to on. By default, the <code>AlertOnDownStates</code> parameter is set as off. When the <code>AlertOnDownStates</code> parameter is set as on, alerts are generated for the VM states configured in both the <code>AbnormalStates</code> and the <code>DownStates</code> parameter list.</p> <div> <p>Note: If HPE Cloud Optimizer is integrated with HPE Operations Manager (HPOM) or HPE Operations Manager <i>i</i> (HPE OMi), you can configure the <code>AlertOnDownStates</code> parameter in the vPV-CustomAlertSensitivityDefinition policy/policy template.</p> </div>
AbnormalStates	<p>By default, the alerts are generated when the VM state changes to any of the VM states configured in the AbnormalStates parameter list.</p> <p>To configure additional states as Abnormal, Down, or Normal, append the state name to the respective parameter list. For example, to include a state as <code>deleted</code> in the Abnormal state list, append <code>deleted</code> to the <code>AbnormalStates</code> parameter list. For example,</p> <pre>AbnormalStates=stuck,crash,hung,unknown,invalid,other,deleted</pre>
DownStates	<p>The alerts are generated when the VM state changes to any of the VM states configured in the DownStates parameter list, when the AlertOnDownStates parameter is set as on.</p>
NormalStates	<p>When the VM changes its state from any Abnormal or Down state to one of the Normal states, the alert is acknowledged in the HPE OM/OMi server.</p>

- After you make the configuration changes, save and close the file.

Managing vCenter Events and Alarms

HPE Cloud Optimizer provides you the capability to collect information about specific vCenter events and alarms. When you integrate HPE Cloud Optimizer with HPE Operations Manager (HPOM) and HPE Operations Manager *i* (OMi), these events and alarms can be forwarded to the HPOM and OMi servers. However, the vCenter alarms and events collected are not shown in HPE Cloud Optimizer alert inbox.

The vCenter events and alarm types for which HPE Cloud Optimizer collects information, is available in the **VIEventTypes.cfg** file.

To configure HPE Cloud Optimizer to forward vCenter Events and Alarms, perform the following steps:

On the HPE Cloud Optimizer machine:

1. Access the `VIEventTypes.cfg` from the following location on the machine:

`/opt/OV/newconfig/OVPM/smepack/VCENTER/data/VIEventTypes.cfg`

The `VIEventTypes.cfg` file contains a list of vCenter event and alarm types collected by HPE Cloud Optimizer. By default, the following event and alarm types are collected by HPE Cloud Optimizer:

```
# Default event filters
VmSuspendedEvent:VmResumingEvent
VmPoweredOffEvent:VmPoweredOnEvent
DrsEnteredStandbyModeEvent:DrsExitedStandbyModeEvent
DrsDisabledEvent:DrsEnabledEvent
VmRenamedEvent
VmRemovedEvent
DrsVmPoweredOnEvent
DrsVmMigratedEvent
NotEnoughResourcesToStartVmEvent
VmBeingHotMigratedEvent
VmFailedMigrateEvent
VmMigratedEvent
VmDiskFailedEvent
VmFailoverFailed
VmNoNetworkAccessEvent
VmUuidChangedEvent
VmUuidConflictEvent
VmOrphanedEvent
HostRemovedEvent
HostShutdownEvent
```


The critical events and their corrective events are written in the same line, separated by a colon. For example,

```
VmSuspendedEvent:VmResumingEvent  
VmPoweredOffEvent:VmPoweredOnEvent  
DrsEnteredStandbyModeEvent:DrsExitedStandbyModeEvent  
DrsDisabledEvent:DrsEnabledEvent
```

To stop the collection for specific alarm or event type, you can comment the specific alarm or event by appending a hash '#' with the alarm or event name. You can add additional events or alarm types by including the type name in the `VIEventTypes.cfg` file. You can uncomment the ones already listed in the `VIEventTypes.cfg` file or include any new vCenter Event or Alarm. To view the complete list of vCenter Events and Alarms, see the *Event Data Object Types* in *VMware documentation*.

Note: If the `VIEventTypes.cfg` file is blank, does not exist, or has all the lines commented out, all the vCenter alarm and event types are considered by HPE Cloud Optimizer for collection.

2. By default, HPE Cloud Optimizer does not collect alarms from the vCenter. To start the collection for alarms, un-comment the following event types by removing the hash '#' preceding the name:

```
# AlarmCreatedEvent  
# AlarmReconfiguredEvent  
# AlarmRemovedEvent  
# AlarmStatusChangedEvent
```

3. Restart HPE Cloud Optimizer using the following command:

```
pv restart
```

On the HPOM/OMi machine:

1. In the **vPV-EventMonitor** policy, set the **AlarmFlag** parameter to **True**. By default, the **AlarmFlag** parameter is to **False** and HPE Cloud Optimizer does not collect alarms from the vCenter.
2. Assign and deploy the **vPV Alerts** aspect to the HPE Cloud Optimizer node. For more information about deploying the **vPV Management Pack** Aspects, see [Deploying the HPE Cloud Optimizer OMi Integration Aspects](#).

HPE Cloud Optimizer can now forward configured events and alarm-related events to HPOM/OMi.

Chapter 8: Capacity

Capacity feature helps you to tune your virtualized infrastructure to maximize efficiency and resource utilization. Using Capacity, you can view the following:

- General overview of the virtualization infrastructure- You can identify the over-utilized and under-utilized resources, idle VMs, and resolve incorrect VM sizing problems.
- Sizing recommendation- Based on the historical resource usage and demand trends and available capacity, HPE Cloud Optimizer recommends ways to optimize your current resources without impacting the service levels or performance of the product.
- Usage overview of entities- You can view the usage summary pages for each entity in the virtualization environment.

Personas

Typical users of this feature are:

- Capacity Planner
- IT Administrator

Navigation

To access the features, select **Capacity**. From the **Capacity** drop-down, select **Overview**, **Forecast**, **Modeler**, or **Placement**.

Using Capacity

You can use the Capacity page to perform the following tasks:

Task	Feature
<ul style="list-style-type: none">• Analyze the capacity, usage and allocation trends for various resources in the virtualized environment• Draw reports based on the utilization and consumption of the resources and focus primarily on reclaiming and adding the resources• Tune your infrastructure for maximum efficiency and reclaiming unused resources	Overview
<ul style="list-style-type: none">• Projection for 30, 60, 90 days and including days to capacity information	Forecast

Task	Feature
<ul style="list-style-type: none">• Determine the use of underutilized resources	
<ul style="list-style-type: none">• Plan the allocation of resources in your environment• Determine the impact of adding resources	Modeler
<ul style="list-style-type: none">• Obtain recommendations to uphold the health of the hosts and datastores• Obtain right-sizing recommendations to resolve incorrect VM sizing problems and achieve better VM performance	Placement

Overview

The Overview feature provides a cohesive view of the resource capacity, usage, allocation, remaining capacity, and opportunity to use the resources more effectively.

Navigation

To access the feature, from the HPE Cloud Optimizer home page, select **Capacity > Overview**. The Optimization Overview page opens.

The following section provides information on the user interface elements on the page.

Graphical Representation

You can view the data represented graphically as Horizontal Bar charts and Pie charts.

Horizontal Bar charts

The horizontal bar charts on the page show the allocation, usage, and capacity details corresponding to the CPU, Memory, and Storage of the resources. For more conceptual information, see ["Appendix" on page 336](#).

The charts available are as follows:

- **CPU (GHz)**- Shows the overall CPU allocation, usage, and capacity of the resources.
- **Memory (GB)**- Shows the overall Memory allocation, usage, and capacity of the resources.
- **Datastore (GB)**- Shows the overall Storage allocation, usage, and capacity of the resources.

Note: When you hover the mouse over a bar, a pop-up shows the actual allocation, usage, or capacity, as applicable.

Pie charts


The pie charts show the VM Activity and sizing data for the VM CPUs and VM Memory.

- **VM CPU Sizing** - Shows the sizing data for the VM CPUs in the datacenter. You can use the legend to identify the number of VMs that are undersized, right sized, and oversized by CPU resources. When you hover the mouse over an area, a pop up shows the number of VMs belonging to the particular sizing category.
- **VM Memory Sizing** - Shows the sizing data for the VM Memory. You can use the legend to identify the number of VMs that are undersized, right-sized, and oversized by memory resources. When you hover the mouse over an area, a pop up shows the number of VMs belonging to the particular sizing category.

You can use the pie chart legend to view the sizing data for the VMs. Click the legend items - **Under Sized**, **Right Sized**, and **Over Sized** to generate the chart based on your requirement. When you click a legend item, the colored circle turns hollow indicating that the item is not selected and corresponding data is removed from the chart. The chart refreshes to display only the solid colored legend items.

Example Scenario: You want to generate and view a chart showing only the undersized and oversized VM CPU data.

Click **Right Sized**. The solid colored circle changes to hollow. The chart refreshes to display only the undersized and oversized VM CPU sizing data.



- **VM Activity** - Shows the data corresponding to the idle and active VMs in the domain. You can use the legend to identify the number of VMs that are idle and active. When you hover the mouse over an area, a pop up shows the number of VMs belonging to the particular category.

Domain Information

You can view the domain information from the table on the right side of the page. The table displays the following information:

- Virtualization Domain: Domain for which the data is shown. For example, vSphere.
- Number of deployed VMs and templates
- Number of powered on VMs in the domain
- Number of remaining VMs in the domain

- Number of Snapshots
- Actions: You can view the forecast, or CPU and memory reservations for the particular VM.

Note: When you drill down and view the **Summary** page for Clusters, Hosts, and VMs, the table also displays the name of the resource(s) up in the hierarchy.

Resources Information Table

The table at the bottom of the page shows the detailed data corresponding to the resources in the domain. The different tabs available are as follows:

- Host
- Datastores

In the Datastores tab, you can filter and view the datastores categorized based on their device latency values. The device latency is the average amount of time, in milliseconds, that a datastore takes to complete a SCSI command.

By default, a datastore is categorized as low performant, if its device latency score is above 20 ms. The time for which the datastore device latency should be above 20 ms, to be categorized as high latency, depends on the configured sensitivity level.

Following are the details for the different sensitivity levels:

High: The Datastore is considered High Latency where the device latency is above 20 ms for 10 percent of the time in a day.

Medium (Default): The Datastore is considered Medium Latency where the device latency is above 20 ms for 20 percent of the time in a day.

Low: The Datastore is considered Low Latency where the device latency is above 20 ms for 30 percent of the time in a day.

For example,

For a Medium Latency scenario, if for 6 hours in a day, which is 25 percent, the device latency of a datastore is greater than 20 ms, we will categorize the datastore as low performant.

Perform the following steps to configure XPL settings for Datastore exclusion:

- a. Log on to HPE Cloud Optimizer system as root user.
- b. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

In this instance,

<install_dir> is the directory, where HPE Cloud Optimizer is installed. The default location is /opt/OV.

- c. A text file opens.
- d. In the text file, create the following namespace:

[pvcd.DatastoreExclusion]

- e. In the [pvcd.DatastoreExclusion] namespace, type the following:

[pvcd.DatastoreExclusion]

IGNORE_PERCENTAGE=<Ignore_Percentage>

LATENCY_THRESHOLD=<Latency_Threshold>


In the above namespace, <Ignore_Percentage> can be an integer between 0 and 100. The default value is 20. This is the percentage of time in a day the Datastore device latency should be above the latency threshold to be categorized as low performant.

<Latency_Threshold> is the threshold value of Datastore device latency. It should be specified in milliseconds (ms). The default value is 20 ms. The threshold value is used to decide if the Datastore is low performant or normal performant.

- Clusters
- Datacenters

Note: The availability of the tabs differ based on the drill down level. For example, if you are viewing the **Summary** page for Cluster, the Datacenters tab is not available.

When you click a tab, the table refreshes to show the data corresponding to the selected resource. You can click the resource attribute name to view the data in ascending or

descending order. You can also click  to print the page.

Searching Data

You can use the search box provided at the top-right corner of the table to search for resources quickly. You can type the required search item and the table refreshes showing the filtered data. For example, if you want to view data corresponding to the clusters, containing the name **Tower**, click **Clusters** and type **Tower** in the search box. The table refreshes to show data corresponding to only the clusters with names containing **Tower**.

The following table provides information about the items shown in the table:

Attribute	Description
Resource Name	<p>Shows the name of the resources in the selected category. For example, if you select Clusters from the tab, this column shows the names of the clusters in the domain.</p> <p>When you click a resource name, the page refreshes to show the summary page for the selected resource.</p>

Attribute	Description
Remaining Capacity	<p>Number of VMs that can be added or created in the Datacenter or Cluster, in addition to the existing ones.</p> <p>The datastores with high I/O latency are not considered while calculating the Reference VMs. View the Summary for Datastore page to view the details of the datastore.</p>
CPU Available for Allocation (GHz)	<p>Shows the total CPU available, in GHz, corresponding to the resource.</p> <p>The value is calculated based on the following formula:</p> $\text{Total Available} = \{((\text{Current Allocation}) * \{(100 - \text{Head Room}\%^1) - \text{Current Util}\}\%^2) / \text{Current Util}\}\%$
CPU Usage (GHz)	Shows the CPU usage, in GHz, corresponding to the resource.
CPU Reclaimable (GHz)	<p>Shows the amount of CPU that can be reclaimed, in GHz, corresponding to the resource.</p> <p>When you click the attribute value, the page refreshes to display the CPU Sizing page for the particular resource. For more information, see CPU Sizing.</p>
Memory Available for Allocation (GB)	<p>Shows the total Memory available, in GB, corresponding to the resource.</p> <p>The value is calculated based on the following formula:</p> $\text{Total Available} = \{((\text{Current Allocation}) * \{(100 - \text{Head Room}\%^3) - \text{Current Util}\}\%^4) / \text{Current Util}\}\%$
Memory Usage (GB)	Shows the Memory usage, in GB, corresponding to the resource.
Memory Reclaimable	Shows the amount of memory that can be reclaimed, in GB, corresponding to the resource.

¹By default, this value is considered as 20%. This means that HPE Cloud Optimizer considers 80% of CPU Util% as healthy.

²CPU Cycles Total Used / Total CPU

³By default, this value is considered as 20%. This means that HPE Cloud Optimizer considers 80% of Memory Util% as healthy.

⁴MemPhysUtil / TotalMem

Attribute	Description
(GB)	When you click the attribute value, the page refreshes to display the Memory Sizing page. For more information, see Memory Sizing .

Summary for Datacenter

Summary for Datacenter page shows an overview of the health and performance of a datacenter in the environment. You can also drill down and view data for the next level of resources in the datacenter. Using the data on this page, you can view the current utilization and plan and allocate your resources better.

Navigation

To view the summary data for a particular datacenter, from the Treemap, right click the datacenter and select **Capacity Usage Trends**. The Summary for Datacenter page opens.

Alternatively, you can also access the page from the Optimization **Overview** page. Click the **Datacenters** tab from the table at the bottom of the page, and select the required datacenter name. The Summary for Datacenter page opens.

Note: To view the summary of all the datacenters, go to the Overview page.

The bar charts provide information on the allocation, usage, and capacity of the CPU, Memory, and DataStore corresponding to the datacenter. The pie charts represent the VM CPU and VM Memory Sizing information and the status of the VMs. For more information, see ["Graphical Representation" on page 171](#).

The tabular data on the right side of the page displays the domain and a count of the clusters, hosts, powered on VMs, and datastores in the particular datacenter. To view more information on each of these resources, you can view the tabular data at the bottom of the page. The resources are grouped and available in tabs. For more information on the table, see ["Resources Information Table" on page 173](#).

The Resource Information table displays data corresponding to clusters, hosts, and datastores for the particular datacenter. You can click the resource names and drill down to view the next level of detailed data.

Viewing Snapshots List

The snapshots information is another aspect to add with the HPE Cloud Optimizer's strong capacity planning and management capabilities. Snapshot captures the state of a VM at a particular time. Snapshots serve as a way for a VM user to roll back in time to a

previous state. However, old snapshots occupy disk capacity unnecessarily. In the Snapshot list page, you can view the list of snapshots and take action to delete old snapshots to reclaim the disk space. The Snapshot information is available at VM, Host, Cluster, and Datacenter levels.

Note: The Snapshot list is available for VMware vSphere environments only.

Navigation

To access the Snapshot list:

1. Select **Capacity > Overview**.
The Overview page opens.
2. In the domain information table, click the link in the No. of Snapshot(s) row.
The Snapshot list page opens.

You can also access the Snapshot list page as follows:

1. Select **Capacity > Overview**.
The Overview page opens.
2. In the resource information table, click either **Host**, **Cluster**, or **Datacenter** tab. The corresponding table containing a list of Host Name, Cluster Name, or Datacenter Name opens. Click a Host Name, Cluster Name, or Datacenter Name, the respective table with Hosts, Clusters or Datacenter opens.
3. Click the Host, Cluster, or Datacenter name. The respective summary page opens. The domain information table lists the number of snapshots.
4. In the No. of Snapshot(s) row, click the link to view the Snapshot list page for that resource.

Resource Information Table

The resource information table in the Snapshot list page contains the following resource information:

- **Name** - Name of the Snapshot.
- **Description** - Description of the snapshot.
- **VM Name** - Name of the VM where the snapshots is present.
- **Created on** - The date when the snapshot was created.

Click a resource attribute name to view the information in ascending or descending order.

Use the search box provided at the top-right corner of the table to search for a resource quickly. Type the required search item in the search box and the table refreshes to show the filtered data.

Summary for Cluster

The Summary for Cluster page shows an overview of data corresponding to a cluster in the environment.

Navigation

To view the summary data for a particular cluster, from the Treemap, right click the cluster and select **Capacity Usage Trends**. The Summary page opens.

Alternatively, you can also access the page from the Optimization **Overview** page. Click the **Clusters** tab from the table at the bottom of the page, and select the required cluster name. The page refreshes to display the Summary for Cluster page.

In this page, you can view the overall health and performance of the cluster. The bar charts provide information on the allocation, usage, and capacity of the CPU, Memory, and Storage corresponding to the cluster resources. The pie charts represent the VM CPU and VM Memory Sizing information and the status of the VMs. For more information, see ["Graphical Representation" on page 171](#).

The tabular data on the right side of the page displays the domain name, datacenter name, and a count of the hosts, VMs, and datastores in the particular datacenter. The table also displays if the cluster is High Availability (HA) enabled and Distributed Resource Scheduler (DRS) enabled. Click on the datacenter name to navigate to the ["Summary for Datacenter" on page 176](#).

To view more information on each of these resources, you can view the tabular data at the bottom of the page. The resources are grouped and available in tabs. For more information on the table, see ["Resources Information Table" on page 173](#).

Summary for Host

The Summary for Host page shows an overview of data corresponding to a host in the environment.


Navigation

To view the summary data for a particular host, from the Treemap, right click the host and select **Capacity Usage Trends**. The Summary page opens.

Alternatively, you can also access the page from the Optimization **Overview** page. Click the **Hosts** tab from the table at the bottom of the page, and select the required cluster name. The page refreshes to display the Summary for Cluster page.


Graphical Representation

In this page, you can view the overall health and performance of the host. The bar charts provide information on the allocation, usage, and capacity of the CPU, Memory, and datastore corresponding to the VMs of the host. The **CPU Utilization** and **Memory Utilization** graphs provide data on the CPU and Memory utilization of the VMs on the host. The legend at the top of the graphs show the VM name and the corresponding color with which it is represented in the graph.



In the preceding example, CPU Utilization graph, the horizontal (X axis) shows the time and the vertical (Y axis) shows the CPU Utilization in GHz, for the host.

When you hover the mouse over a colored area, a pop up appears showing the VM name and CPU utilization, in GHz, specific to the VM. To view more information on the VM, click the corresponding colored area. The graph refreshes to display data for the particular VM. The following figure shows the CPU Utilization graph for a particular VM on the host.



In the preceding figure, the graph shows the CPU utilization for the selected VM. The vertical axis shows the utilization specific to the VM in GHz.

To go back and view the previous graph for the host, click on the graph area.

Domain Information

The domain information table on the right side of the page displays more details about the host. The following table lists the details.

Item Name	Description
Virtualization	The virtualization domain to which the host belongs.

Item Name	Description
Domain	
Cluster Name	Name of the cluster to which the host belongs. Click the cluster name to view the "Summary for Cluster" on page 178 .
Datacenter Name	Name of the datacenter, of which, the host is a part. Click on the datacenter name to view "Summary for Datacenter" on page 176 .
vCenter Name	Displays the name of the vCenter to which the host belongs.
Powered On VMs	Displays the count of VMs on the host, that are in the powered on state.
Logical CPUs	Displays the number of logical CPUs the VM has.
CPU Reservation (GHz) (VM Kernel)	Indicates the CPU Reservation for the host. Processes and services running on the host make CPU Reservation requests to ensure a minimum guaranteed execution rate. CPU reservations are usually in the form of x units of time out of every y units for a process.
CPU Shares Priority	Displays the CPU shares percentage of the VM. CPU share is a portion of the CPU resources of the system, that is allocated for a process. When creating a VM, a CPU share priority and a number of CPU cores is allocated.
CPU Clock Speed (GHz)	Shows the clock speed of the CPU, in GHz. Clock is a microchip regulating the timing and speed of system functions. To execute an instruction, the CPU requires a particular number of clock cycles. A system performs faster when the CPU speed is also faster.
Memory Reservation (GB) (VM Kernel)	Displays the memory reserved for the activities running on the host.

Resource Information Table

The table lists the VMs and datastores for the host under the **VMs** and **DataStores** tabs, respectively. The following table lists the details available under the **VMs** tab.

Note: The VMs tab is available only when you are viewing the Summary for Host page.

Item Name	Description
VM Name	Name of the VM.
Recommended Logical CPUs Number	Number of logical CPUs that are recommended for the VM to perform better.
CPU Reclaimable (GHz)	Shows the amount of CPU that can be reclaimed.
Recommended Memory (GB)	Recommended memory that must be allocated to programs and services on the VM, for the VM to perform better.
Memory Reclaimable (GB)	Shows the amount of memory that can be reclaimed.

Summary for VM

The Summary for VM page shows an overview of data corresponding to a VM in the environment.

Navigation

To view the summary data for a particular VM, from the Treemap, right click the VM and select **Capacity Usage Trends**. The Summary page opens.

Alternatively, you can also access the page from the Optimization **Overview** page. Click **Hosts** tab and select the host to which the particular VM belongs. The Summary for Host page appears. Click **VMs** tab, select the required VM from the table at the bottom of the page, and select the required VM name. The page refreshes to display the Summary for VM page.

The page provides a detailed overview of the utilization and performance of the VM.

Configured

Shows the configured values for the following attributes of the VM:

- Logical CPUs
- Memory Configuration (GB)

- Memory Reservation (GB)
- CPU Reservation (GHz)

Recommended

Shows the recommended values for the following attributes of the VM:

- Logical CPUs
- Memory Configuration (GB)
- Memory Reservation (GB)
- CPU Reservation (GHz)

Remarks

Provides information on the overall status of the CPU and Memory of the VM.

Graphical Representation

Shows the performance of the VM graphically. The graph legends help you identify the colors associated with each item in the graph. To view data corresponding to a particular property of the VM, you can click the legend items and the graph refreshes.

The graphs available are as follows:

- **CPU and Memory Utilization** - Shows the CPU Utilization and Memory Utilization of the VM.
- **SecondChartHeader** - Shows the following data:
 - CPU Co-Stop Time (in milliseconds)
 - CPU Ready Time (in milliseconds)
 - CPU Demand (in MHz)
- **Physical CPU and Memory Utilization** - Shows the following data:
 - VM Physical Memory Utilization
 - VM Physical CPU Utilization
- **Latency** - Shows the following data:
 - Disk Read Latency
 - Disk Write Latency

Domain Information Table

The table lists the following information.

Item	Description
Virtualization Domain	Name of the virtualization domain to which the VM belongs.
IP Address	IP Address of the VM.
VM Hosted On	Name of the host to which the VM belongs. Click the host name to navigate to the "Summary for Host" on page 178 .
Cluster Name	Name of the cluster, of which, the VM is a part. Click the cluster name to navigate to the "Summary for Cluster" on page 178 .
Datacenter Name	Name of the datacenter to which the VM belongs. Click the datacenter name to navigate to the "Summary for Datacenter" on page 176 .
vCenter Name	Name of the vCenter to which the VM belongs.
First Data point	Shows the time interval when the data was logged first.
Last Data Point	Shows the time interval when the data was logged last.
CPU Cycles Used (GHz)	Shows the CPU cycles used to process the instructions. Clock is a microchip regulating the timing and speed of system functions. Each clock tick is a clock cycles. To execute an instruction, the CPU requires a particular number of clock cycles.
CPU Clock Speed (GHz)	Shows the clock speed of the CPU, in GHz. A system performs faster when the CPU speed is also faster. CPU speed determines how many instructions it can perform in a second.
CPU Limit (GHz)	Shows the limit to which the CPU can be used for the VM activities. CPU limits are set to limit the CPU consumption from going high on VMs. Setting CPU limits help in managing contention issues that may arise in your environment.
CPU Reclaimable (GHz)	Shows the amount of CPU that can be reclaimed.

Item	Description
Memory Shares	Displays the memory shares of the VM. Memory share is a portion of the memory of the system that is allocated for a process. When creating a VM, memory shares are also allocated.
Mem Used (GB)	Shows the memory used by the VM activities.
Memory Reclaimable (GB)	Shows the amount of memory that can be reclaimed.
Datastore Usage (GB)	Lists the datastores corresponding to the VM. Click the datastore names to navigate to the "Summary for Datastore" below .

Summary for Datastore

The Summary for Datastore page shows an overview of data corresponding to a Datastore in the environment.

Navigation

You can access the Datastore Summary in the following ways:

- From the Treemap
In the Treemap, right-click the Datastore and select **Capacity Usage Trends**. The Datastore Summary page appears.
- From the Capacity Overview page
 1. Go to **Capacity > Overview**.
 2. From the table at the bottom of the page, click the **Datastores** tab.
 3. In the Datastore Name column, click the Datastore which you want to view the Summary. The Datastore Summary page appears.

Graphical Representation

The bar chart shows the storage allocation, usage, and capacity of the Datastore.

The pie charts available are as follows:

- Utilization by File Types (GB)
- Top 5 VM's Disk Usage (GB)
- Activity Status for VMs

Remarks

Provides information on the overall space and health of the Datastore.

Domain Information

The table at the side lists the following information:

- Virtualization Domain
- vCenter Name
- Datacenter Name
- Datastore Type
- I/O Size Analysis
- VMs Avg. Disk Read Latency (ms)
- VMs Avg. Disk Write Latency (ms)
- Usage for Active VMs (GB)
- Reclaimable Orphan VDisks space (GB)

Resource Information Table

The table lists the VMs, Hosts, and vDisks associated with the respective Datastore.

VMs

You can see the following information for the VMs:

- VM Name
- Activity state of the VM
- Space Provisioned in GB - This value is calculated based on only the Active VMs.
- Disk Usage in GB - This is the disk usage in GB for a VM.
- Disk Read Latency (ms)
- Disk Write Latency (ms)

vDisks

Virtual machines use virtual disks (vDisks) for their operating system, application software, and other data files. A virtual disk is stored as a VMDK file on a Datastore. Orphaned vDisks are not connected to any virtual machine.

Orphaned vDisks get created when the VMs are deleted from the inventory but the associated vDisks are not deleted. The deleted VM's CPU and Memory is released,

however, the storage remains unavailable for utilization. Also, when migrating VMs from one host to another host, if migration process fails due to some reason, vDisks get created in the source host.

Hence it is important to identify Orphaned vDisks that can be deleted and the storage to be released for utilization. Any vDisk modification will reflect after four hours in HPE Cloud Optimizer.

You can view the following information for vDisks:

- vDisk Name
- Folder Name
- Status - Powered On, Powered Off, Suspended, or Orphaned
- Disk Usage (GB)
- VM Name

Hosts

You can view the following information for Hosts:

- Host Name
- Disk Write Latency (ms)
- Disk Read Latency (ms)

CPU Sizing

The CPU Sizing page shows the CPU sizing data for the VMs, based on the historical resource usage. You can view the sizing data for all the VMs corresponding to a resource. The pie chart on the page graphically represents the number of the VMs that are Right sized, Over sized, and Under sized. Hovering the mouse over the shaded area shows the number of VMs in the particular category.

The page also contains a search box to filter data based on your requirement. For more information, see ["Searching Data" on page 174](#)

Navigation

You can access the CPU Sizing page from the following:

- Resource Information table available on the **Overview** and Summary pages - When you click the CPU Reclaimable value corresponding to a resource, the page refreshes to show the CPU sizing data of the VMs for the resource. The following table lists the tabs in the Resource Information table from which you can access the CPU Sizing page.

Tab Name	Page
Datacenter	CPU Sizing for Datacenter
Cluster	CPU Sizing for Cluster
Host	CPU Sizing for Host

- Toolbar - Click **Reports** from the Toolbar on top of the page, and select **VMs CPU Sizing** from the drop down.

Resource Information Table

The table contains the following four tabs:

- **All** - Lists all the VMs corresponding to the resource.
- **Right Sized** - Lists the VMs that are Right sized.
- **Over Sized** - Lists the VMs that are Over sized.
- **Under Sized** - Lists the VMs that are Under sized.

All the tabs contain the following data corresponding to the VMs:

- **VM Name** - Name of the VM. You can click the VM name to navigate to the ["Summary for VM" on page 181](#).
- **VM Sizing** - Shows whether the VM is Right sized, Under sized, or Over sized.
- **Current CPU** - Shows the current CPU allocated for the VM.
- **Recommended CPU** - Shows the CPU that must be assigned to the VM to make it Right sized.

For the VMs listed under the Right Sized tab, the value of Current CPU and Recommended CPU is the same.

- **Actions** - Following actions can be taken for the particular VM:
 - **Capacity Usage Trends**: Manage the resource utilization in your environment.
 - **Workbench**: Draw graphs and generate reports for performance data.
 - **VM Forecast**: View forecast of the resources utilized for a particular VM.

With the help of CPU sizing details, you can plan to consume the Right sized, Over sized, and Under sized resources effectively. You can use the Modeler feature to analyze and allocate the Over sized and Under sized resources based on the workloads. Using this feature, HPE Cloud Optimizer gives you an overview of the CPU, Memory, and Storage Utilization of the resources in your environment.

Memory Sizing

The Memory Sizing page shows the Memory sizing data for the VMs. You can view the sizing data for all the VMs corresponding to a resource. The pie chart on the page graphically represents the number of the VMs that are Right sized, Over sized, and Under sized. Hovering the mouse over the shaded area shows the number of VMs in the particular category.

The page also contains a search box to filter data based on your requirement. For more information, see ["Searching Data" on page 174](#)

Navigation

You can access the Memory Sizing page from the following:

- Resource Information table available on the **Overview** and Summary pages - When you click the Memory Reclaimable value corresponding to a resource, the page refreshes to show the Memory sizing data of the VMs for the resource. The following table lists the tabs in the Resource Information table of the **Overview** page, from which you can access the Memory Sizing page.

Tab Name	Page
Datacenter	Memory Sizing for Datacenter
Cluster	Memory Sizing for Cluster
Host	Memory Sizing for Host

- Toolbar - Click **Reports** from the Toolbar on top of the page, and select **VMs Memory Sizing** from the drop down.

Resource Information Table

The table contains the following four tabs:

- **All** - Lists all the VMs corresponding to the resource.
- **Right Sized** - Lists the VMs that are Right sized.
- **Over Sized** - Lists the VMs that are Over sized.
- **Under Sized** - Lists the VMs that are Under sized.

All the tabs contain the following data corresponding to the VMs:

- **VM Name** - Name of the VM. You can click the VM name to navigate to the ["Summary for VM" on page 181](#).
- **VM Sizing** - Shows whether the VM is Right sized, Under sized, or Over sized.
- **Current Memory** - Shows the current Memory allocated to the VM.

- **Recommended Memory** - Shows the Memory that must be assigned to the VM to make it Right sized.

For the VMs listed under the Right Sized tab, the value of Current Memory and Recommended Memory is the same.

- **Actions** - Following actions can be taken for the particular VM:
 - **Capacity Usage Trends:** Manage the resource utilization in your environment.
 - **Workbench:** Draw graphs and generate reports for performance data.
 - **VM Forecast:** View forecast of the resources utilized for a particular VM.

With the help of Memory sizing details, you can plan to consume the Right sized, Over sized, and Under sized resources effectively. You can use the Modeler feature to analyze and allocate the Over sized and Under sized resources based on the workloads. Using this feature, HPE Cloud Optimizer gives you an overview of the CPU, Memory, and Storage Utilization of the resources in your environment.

Configuring Memory Sizing

Memory sizing is the planning for the memory calculation that helps to tune your virtualized environment effectively. This makes the environment more efficient. You can configure memory sizing by selecting the metrics used to calculate memory sizing.

For instance, you can configure memory sizing to be based on consumed memory per VM. The default metric is based on active memory usage. This, however, needs to be done at a business group level.

Create a business group of the VMs to configure memory sizing. Name, Annotation, IP Address, or Host Name is used to search the VMs to create a business group. Use keyword or regular search expressions to search for VMs.

See the [creating a business group](#) section for detailed information to create a business group.

Note: Make sure that a VM is not part of multiple business groups.

After creating the business group, follow these steps to configure the memory sizing metric:

1. Log on to the HPE Cloud Optimizer system as a root user.
2. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

In this instance, <install_dir> is the directory, where HPE Cloud Optimizer is installed. The default location is /opt/OV.

3. A text file opens.
4. In the text file, create the following namespace:

[pvcd.po.grouping.<grouping name>]

<group name> = metric or expression

Example:

DBGroup = MemUsed

Follow these steps to change the default metric:

1. Log on to the HPE Cloud Optimizer system as a root user.
2. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

In this instance, <install_dir> is the directory, where HPE Cloud Optimizer is installed. The default location is /opt/OV.

3. A text file opens.
4. In the text file, create the following namespace:

[pvcd.PO]

DEFAULT_MEM_SIZING_METRIC = <MemActive>

<MemActive> - The name of the memory related metric.

```
[pvcd.PO]
ANALYSIS_END_DATE=1449650304
ANALYSIS_START_DATE=1449628499
COLLINTRLSECS=300
DBUPDATE_FREQUENCY=360
DEFAULT_MEM_SIZING_METRIC=MemActive
INSTANCEALIVEWINDOWMINS=11
OVERSUBSCRIPTION_CPU=100
OVERSUBSCRIPTION_MEM=100
REFERENCE_VM_CPU=2.0
REFERENCE_VM_DISK=20.0
REFERENCE_VM_MEM=4.0
RIGHTSIZE_CPU_HEADROOM=20
RIGHTSIZE_DISK_HEADROOM=10
RIGHTSIZE_MEM_HEADROOM=20
SIZINGTYPE=2
```

View Idle VMs

The View Idle VMs page shows the data for the VMs in idle state. You can view the memory and CPU utilization for all the idle VMs.

The page also contains a search box to filter data based on your requirement. For more information, see ["Searching Data" on page 174](#)

Navigation

To access the Idle VM list:

1. Select **Capacity > Overview**.
The Overview page appears.
2. In the VM Activity section, click the **View Idle VMs** link.

Resource Information Table

The table contains the following data corresponding to the idle VMs.

- **VM Name** - Name of the VM. You can click on the VM name to navigate to ["Summary for VM" on page 181](#).
- **VM CPU Util %** - Shows the current CPU allocated to the VM in percentage.
- **VM Memory Util %** - Shows the current memory allocated to the VM in percentage.
- **Avg. Disk I/O** - Shows the average disk input/output.
- **Avg. Network I/O** - Shows the average network input/output.
- **Actions** - Following actions can be taken for the particular VM:
 - **Capacity Usage Trends:** Manage the resource utilization in your environment.
 - **Workbench:** Draw graphs and generate reports for performance data.
 - **VM Forecast:** View forecast of the resources utilized for a particular VM.

With the help of Idle VM details, you can plan to consume the underutilized resources effectively. You can use the Modeler feature to analyze and allocate the idle resources based on increasing the workloads. Using this feature, HPE Cloud Optimizer gives you an overview of the CPU, Memory, and Storage Utilization of the resources in your environment.

View Powered off VMs

The View Powered off VMs page shows the data for the VMs in powered off state. You can view the time since when the VMs were powered off, down time, and disk usage.

The page also contains a search box to filter data based on your requirement. For more information, see ["Searching Data" on page 174](#)

Navigation

To access the Powered Off VMs list:

1. Select **Capacity > Overview**.

The Overview page appears.

2. In the VM Activity section, click the **View Powered Off VMs** link.

Resource Information Table

The table contains the following data corresponding to the powered off VMs.

- **VM Name** - The name of the VM. You can click on the VM name to navigate to ["Summary for VM" on page 181](#).
- **Powered off Since:** - Shows the time from when the VM is powered off.
- **Down Time %** - Shows the percentage of time for which the VM is down.
- **Host Name** - The name of the host.
- **Disk Usage (GB)** - The amount of disk, in GB, used.

With the help of powered off VM details, you can plan resources effectively. You can use the Modeler feature to analyze and allocate the powered off resources based on increasing the workloads. Using this feature, HPE Cloud Optimizer gives you an overview of the CPU, Memory, and Storage Utilization of the resources in your environment.

CPU and Memory Reservation

This page provides an overview of the CPU and Memory reservations corresponding to VMs in the environment.

Memory Reservation is the process by which a certain amount of memory is set to be always accessible for the VM. The value is shown in GB.

CPU Reservation is the process by which a certain amount of the CPU process is set to be always accessible for the VM processes. The value is shown in GHz.

You can also navigate to the corresponding Summary pages of the VMs from here. A search box is provided to filter and view data, based on your requirement. For more information, see ["Searching Data" on page 174](#).

Navigation

You can navigate to this page from the toolbar on top of the page. Click **Reports** and select **VMs sorted by Reservation**.

Resource Information Table

The table lists the following data:

- **Name** - Name of the VM. You can click the VM name to navigate to the ["Summary for VM" on page 181](#).

- **Cluster Name** - Name of the cluster to which the VM belongs.
- **DC Name** - Name of the Datacenter to which the VM belongs.
- **CPU Reservation** - Shows the CPU reserved for the VM activities, in GHz.
- **Memory Reservation** - Shows the memory reserved for the VM activities, in GB.

Use Case: Using Capacity Usage Trends and Optimization Option to Manage the Resource Utilization in your Environment

The following section shows you a scenario where you can find out the reason for the resource utilization, check the capacity usage trend, and then allocate the resources accordingly.

Scenario

John is a VMware Administrator and is monitoring the performance of all the entities in the Virtualization environment. He wants to use the resources optimally, thereby ensuring reliable performance at an optimal cost.



He observes that the usage of one of the Datacenter is close to 70% and wants to add more resources to this datacenter. He wants to check the utilization and optimize the environment. Also, wants to analyze where he can reclaim the under utilized resources or add the resources as required.

The steps include the following:


1. Log on to HPE Cloud Optimizer user interface.
2. Go to the **Treemap view > Datacenter** to view the overall usage.
3. Select the Datacenter that is showing red or orange color. Right-click on the selected Datacenter and select **Capacity Usage Trends** option. The **Overview** page opens. Alternatively, from the HPE Cloud Optimizer home page, you can open the **Overview** page which contains the details of the datacenter, cluster, hosts, and datastores.

Points that you can gather from the Overview page are:

Overall view of the Datacenter. Check the number of cluster, hosts, and powered on VMs available for the selected Datacenter.

Virtualization Domain	vSphere
vCenter Name	16.184.46.174
Powered On VMs	51
No. of Snapshot(s)	0
Actions	 

Additionally, you can view the overall CPU, Memory and Storage Utilization from the Overview page. Here, for example, you can view that the CPU allocation is much higher than the capacity. The usage of CPU utilization is very less. When the usage will increase and become close to the capacity, then you may want to add more resources or check for the reclaimable resources.



From the pie charts, you can derive if the resources are undersized, oversized or right sized. For more information, see ["Graphical Representation" on page 171](#).

You can navigate to get additional details about Cluster, Hosts and Datastore from the Overview page. The attributes available help you to drill down to check the CPU and Memory availability, based on the allocation, usage, Memory and CPU reclaimable.

CPU Attributes for the selected Cluster

Host Name	Remaining Capacity (VMs)	CPU Available for Allocation (GHz)	CPU Usage (GHz)	CPU Reclaimable (GHz)	Memory Available for Allocation (GB)	Memory Usage (GB)	Memory Reclaimable (GB)
-----------	--------------------------	------------------------------------	-----------------	-----------------------	--------------------------------------	-------------------	-------------------------

If you view that CPU Reclaimable for one of the Cluster is 530.53 GHz, then you can navigate to check the right sized, under sized, and over sized CPU utilization. Click the **CPU reclaimable** value and the **CPU Sizing** page opens. You can view the VM name, sizing, allocated CPU and recommended CPU. The pie charts show the over sized, under sized, and right sized VMs. For more information, see ["Graphical Representation" on page 171](#).

You can drill down and check the undersized VM. HPE provides some suggestions such as the Recommended CPU and so on. You can follow the suggestions and allocate more CPUs to this undersized VM.

Under Sized VM Details

All	Right Sized	Over Sized	Under Sized			
VM Name	VM Sizing		Current CPU	Recommended CPU		
RHEL6.2_x64_NMC_pravekum_nmcvm121	Under Sized		1	2		


To find additional information about the VM that is undersized, you can click the VM and can check the VM utilization. The VM Summary page opens.

Points that you can gather from the VM Summary page are:

View the overview about the VM on the right panel. You get the details of the VM such as where the VM is hosted, which data center it belongs to and so on. Suggestions about the resource usage for the selected VM is available in the VM Summary page.

You can get the resource usage trend of a VM and check the following graphs for more details:


View the CPU and Memory Utilization graph to derive and conclude the Allocation versus Usage. You can also validate the recommendation that HPE provides by categorizing the VMs as under sized or over sized.



Conclusions from the Graph:

- Current CPU utilization is high and if you add more resources then the CPU utilization will be optimal.
- Memory utilization is low and you can safely reclaim the allocated memory.

View the CPU and Memory Contention.



You can conclude the following from this example:

- Analyze the resource allocation based on the bar graphs and pie charts.
- Find all the VMs that are available in the Datacenters, Clusters, or Hosts.
- Drill down to the VM where HPE suggestions are available for the selected VM and use the recommendation accordingly.
- Check the resource usage. You can optimize the environment by allocating the resources based on the current capacity trends.

Using Placement

You can use the Placement feature to properly plan and allocate your resources in the virtualized environment. The feature provides you the recommendations on the resources where you can allocate the new VMs.

Note: HPE Cloud Optimizer will not consider the remaining unutilized reserved memory of any VM while giving recommendation for placement.

For example, in a Host with 5 VMs having following configurations:

VMs	Reserved Memory	Configured Memory	Usage	Unutilized Reserved Memory (Reserved Memory - Usage)
VM1	2 GB	4 GB	1 GB	1 GB
VM2	2 GB	4 GB	1.5 GB	0.5 GB
VM3	2 GB	4 GB	3 GB	0
VM4	0	4 GB	0.5 GB	0
VM5	1 GB	4 GB	1 GB	0
			Total	1.5 GB

This 1.5 GB of memory will not be taken into consideration while giving placement recommendation.

Placement Recommendation

After providing the details of the VMs that you intend to add to your environment, HPE Cloud Optimizer provides you the placement recommendation for the new VMs.

These recommendations are based on remaining capacity and performance considerations of the resources. HPE Cloud Optimizer does not include a Host-Datastore combination, whose latency is greater than the configured value.

The feature lists the following:

- Hosts to which you can add the VMs
- Corresponding Cluster and Datacenter to which the host belongs
- Number of VMs that you can add to each host

When you click the datacenter name, cluster name, or host name, you can navigate to the specific Summary page.

Navigation

To access the feature, from the HPE Cloud Optimizer home page, select **Capacity > Placement**. The **New VM Placement** page opens.

Allocate new VMs

To view the recommendation on where to allocate new VMs in your environment, follow these steps:

1. Select **Placement** from the **Capacity** drop-down list in the HPE Cloud Optimizer home page.
The **New VM Placement** window opens.
2. Type the number of VMs you want to add, in the **Number of Virtual Machines** field.
3. Type the CPU size of the VMs in the **VM Number of vCPU** field.
4. Type the Memory size of the VMs in the **VM Memory Size (GB)** field.
5. Type the Disk size of the VMs in the **VM Disk Size (GB)** field.
6. Select the virtualization domain from the **Domain** drop-down list.
7. Click **Add new VM(s)**.

The VM information appears in the **Placement Request List** table.

8. Click **Find Placement**.

The **Placement Recommendation** for the VMs appears on the page.

Note: HPE Cloud Optimizer provides placement suggestions for **vSphere** domain only if all the requested VMs can be placed in the vCenters.

Forecast

Based on historical usage trends, HPE Cloud Optimizer can forecast the resource utilization of entities in your environment. You can use this data to plan for future resource requirements and allocations. HPE Cloud Optimizer uses forecasting techniques to predict the resource utilization based on current allocation.

The Forecast window opens showing the data for the resources. After installing HPE Cloud Optimizer, based on the usage trends of resources, HPE Cloud Optimizer starts forecasting the resource utilization.

You can view and drill down to the forecast data for the following entities:

- Aggregate of all domains
- Datacenter
- Cluster
- Host
- VM
- Datastore

You can drill down to the entities in the datacenter and view the forecast data corresponding to each of the entities. Click the entity names to view the forecast data corresponding to each. The table at the bottom of the page displays the entities to which you can drill down. The value of the resource utilization corresponding to the capacity and days to capacity for all the entities at 30 days, 60 days, and 90 days, respectively appears in the graph.

The Forecast page displays forecast results for entities based on the analysis done in the last planning and optimization analysis cycle. If you do not see a particular entity in the Forecast page that has been added recently or has become active recently, wait for the next planning and optimization analysis cycle to begin or run the analysis immediately.

The details of the planning and optimization cycle can be viewed from **Settings > Collection & Analysis** tab. The analysis runs automatically every six hours. However, if required you can run the analysis at any time.

Note: The minimum requirement to get accurate graphs for forecast is that the data should be collected at least for two cycles. If you have sufficient data available in

database, you get better graphs. Once data is available, the tool shows the derived data for a period of 90 days.

Navigation

You can access the Forecast in the following ways:

- From **Capacity > Forecast**
 1. From the top-left corner, click **Environment**.
 2. Select the required entity from the domain. You can also search for the entity.
- From **Overview > Treemap**
 1. In the Options pane, select the Domain from the Domain drop down. The page refreshes to give the details of the resources of the selected domain.
 2. Right-click the entity and select Forecast. The Forecast page opens.

HPE Cloud Optimizer can forecast the following:

- CPU Usage
- Memory Usage
- Storage Usage

Note: Storage Usage data is not available for VMs and Hosts in KVM/Xen domain.

Data Representation

The forecast data is represented in the form of graphs for ease of use. The data is plotted against dates and usage. The graph shows the Forecast and Metric.

- **Forecast** - Shows the forecast data of the resource utilization for the next 90 days.
- **Actual**-Shows the actual resource utilization data for the last 90 days. The actual usage is the six hourly summarization. This means that every 6 hours, one aggregated (averaged) data point gets added. The forecast calculation happens using every data sample in database and not on 6 hourly aggregates.

Note: If the difference between the Actual and the Forecast value is very less, you may see two different points plotted on the graph showing the same Actual and Forecast value on mouse-over. This is because, on mouse-over, decimal values are truncated up to two decimal places.

For example, if the Actual value is 37.50 and the forecasted value is 37.501, you may see two different points plotted on the graph showing the same value of 37.50 for both the points.

Tabular Representation

- **Capacity** - Shows the maximum value for the resource capacity. For example, for Disk Usage graph, it shows the maximum storage capacity of the entity.
- **Days to Capacity** - Shows the number of days in which the entity reaches the maximum capacity.


Note: Sometimes the forecast calculation shows that the capacity is already reached, but it has not reached capacity in reality. This situation arises if the usage was increasing steadily earlier and suddenly usage growth reduced. In such a case, the Days to Capacity shows as **0** days. If the usage continues to be steady, or if it declines, the forecast calculation gets corrected in due course of time.

- **30 Days, 60 Days, 90 Days** - Shows the resource utilization value at 30 days, 60 days, and 90 days, respectively.
- **Best Fit Curve** - Shows the type of curve the graph represents. The best fit curve for each of CPU, Memory, and Storage of each entity is chosen based on the value of R². The curve with the highest R² is chosen by HPE Cloud Optimizer as the best fit curve. The values can be Linear, Exponential, Logarithmic, or Power.
- **R²** - This is the Coefficient of Determination. The value for this indicates how well the data points fit on a line or curve. The value ranges between 0 and 1. The forecast is more correct when the value is closer to 1. When historic data is very less, based on which forecasting is done, the value is close to 0.



Example:

The aggregated cluster forecast report contains data from the available Hosts and VMs.

Check the following forecast report for the memory utilization:



When you view from the aggregated view that the **Days to Capacity** is 8 days, then check the capacity of the hosts in the specific cluster.

DataStores				
Hosts				
Host Names	CPU		Memory	
	Days To Capacity	Capacity(GHz)	Days To Capacity	Capacity(GB)
 nnmvz31.ind.hp.com	20	27.19	5	142.33
 nnmvz32.ind.hp.com	> 90	27.19	> 90	142.33

One of the host shows days to capacity as 5 days.

Conclusion: Here, the cluster is the container and the entities belonging to this container are the 2 hosts.

Forecast data at the aggregated level (cluster level) is calculated based on the utilization of the current resources for the entities (host 1 and host 2) available in the container. The forecast data (days to capacity) for the host1 and host 2 will not be taken into consideration while calculating the forecast data (days to capacity) of the cluster.


Therefore, forecast data available at the aggregated level is independent of the forecast data available for the entities belonging to the same container.

Forecasting CPU Utilization

The CPU Usage graph displays the CPU utilization forecast for the entities in your environment. When you drill down to the next level entities, you can view the graphs corresponding to each entity in the environment.

Understanding CPU Usage Forecast

The following figure shows a sample CPU Usage forecast graph.



Derivation

Following are the derivations from the preceding figure:


- Blue line shows the forecast data and CPU usage.
- Red line shows the actual CPU usage of the entity.
- When you hover the mouse over an area, a pop up shows the forecast CPU usage at that date. The x-axis shows the xx.xx value and the y axis shows the dates. If the plotting in x axis starts from xx.xxx to xx.xxy, then the value gets truncated and is only displayed as xx.xx to xx.xx. Ignore the values displayed in the x axis. Only check the plotted graph.

Forecasting Memory Utilization

The Memory Usage graph displays the Memory utilization forecast for the entities in your environment. When you drill down to the next level entities, you can view the graphs corresponding to each entity in the environment.

Understanding Memory Usage Forecast

The following figure shows a sample Memory Usage forecast graph.



Derivation

Following are the derivations from the preceding figure:

- Blue line shows the forecast data and Memory usage.
- Red line shows the actual Memory usage of the entity.
- When you hover the mouse over an area, a pop up shows the forecast Memory usage at that date. The x-axis shows the xx.xx value and the y axis shows the dates. If the plotting in x axis starts from xx.xxx to xx.xxy, then the value gets truncated and is only displayed as xx.xx to xx.xx. Ignore the values displayed in the x axis. Only check the plotted graph.


Forecasting Disk Utilization

The Disk Usage graph displays the disk utilization forecast for the entities in your environment. When you drill down to the next level entities, you can view the graphs corresponding to each entity in the environment.

Note: You must set the **Browse Datastore** role available under **Datastore** in VMware vCenter to view the Disk Utilization forecast data for VMs.

Understanding Disk Usage Forecast

The following figure shows a sample Disk Usage forecast graph.



Derivation

Following are the derivations from the preceding figure:

- Blue line shows the forecast Disk Usage data.
- Red line shows the actual Disk usage of the entity.
- When you hover the mouse over an area, a pop up shows the forecast Disk usage at that date. The x-axis shows the xx.xx value and the y axis shows the dates. If the plotting in x axis starts from xx.xxx to xx.xxy, then the value gets truncated and is only displayed as xx.xx to xx.xx. Ignore the values displayed in the x axis. Only check the plotted graph.

Using Forecast to plan for resources

This section shows how you can use the Forecast feature to plan for resources.

Scenario

Mark is a VMware Administrator in an organization. He wants to know the current disk usage of the datacenter and plan for any hardware requirements that may arise. He uses the Forecast feature for this purpose.

The tasks that Mark performs are as follows:

1. Log in to HPE Cloud Optimizer.
2. Right click the required datacenter, from the Treemap, and select **Forecast**.
3. View the graphs and tabular data appearing on the Forecast page.

After analyzing the graphs and tabular data, Mark arrives at the conclusion that the disk usage will reach its maximum capacity in a month.

Actions

Mark performs the following actions based on the forecast data:

- Plans to increase the disk capacity for the datacenter
- Uses the Optimization feature to check the datastore allocation for the entities in the datacenter and reclaim any wasted resources. For more information, see "[Capacity](#)" on page 170

Modeler

In a virtualization environment, there always arises a need to increase resources, based on increasing workloads and other requirements. As an Administrator, before adding resources and entities, you may want to analyze how the allocation of resources impacts your environment. It is important to get accurate results from the analysis. Based on the analysis, you can plan and allocate the resources effectively.

You can use the Modeler feature in HPE Cloud Optimizer to do an analysis of how your environment performs after allocating the resources. After you provide the details of the entities to be added or removed, HPE Cloud Optimizer gives you an overview of how the CPU, Memory, and Storage Utilization is impacted based on the plan. You can use the results from the analysis to determine if you are allocating the resources effectively or need to re-plan the resource allocation to suit your environment needs.

Note: HPE Cloud Optimizer will not consider the remaining unutilized reserved memory of any VM while giving recommendation for CPU and memory.

For example, in a Host with 3 VMs, having following configurations:

VMs	Reserved Memory	Configured Memory	Usage	Unutilized Reserved Memory (Reserved Memory - Usage)
VM1	2 GB	4 GB	1 GB	1 GB
VM2	2 GB	4 GB	1.5 GB	0.5 GB
VM3	2 GB	4 GB	3 GB	0
VM4	0	4 GB	0.5 GB	0
VM5	1 GB	4 GB	1 GB	0
			Total	1.5 GB

This 1.5 GB of memory will not be taken into consideration while giving CPU or memory recommendation.

For example, if you want to add 50 more VMs in your environment, based on a new project requirement, you can use Modeler in HPE Cloud Optimizer and provide the details of the VMs to be added. After the analysis, HPE Cloud Optimizer returns you the information on how the CPU, Memory, and Storage Utilization of your environment is impacted after adding the 50 VMs. If any of these resources are impacted, you can try allocating the VMs across datacenters and clusters and analyze the performance for

each allocation. Once you find a suitable allocation you can go ahead and add the resource in your environment.

Note: You can do capacity modeling only for VMware and KVM domain. You can also do capacity analysis for Business Groups of Standalone Hosts.

Navigation

To access the feature, from the HPE Cloud Optimizer home page, select **Capacity > Modeler**.

Using the Capacity Modeler

Capacity Modeling for VMware

You can do capacity modeling for VMware at the Datacenter or Cluster level.

Capacity Modeling for KVM

To do capacity modeling for KVM, you must create a Business Grouping of KVM Hosts.

For information on how to create Business Groups, see [Business Groups](#).

Capacity Modeling for Standalone Host

To do capacity modeling for Standalone Hosts, you must create a Business Grouping of Standalone Hosts.

For information on how to create Business Groups, see [Business Groups](#).

Domain Information

You can view the domain information from the table on the right side of the page. The table displays the following information:

- Name of the selected Datacenter or Cluster
- Number of remaining VMs. You can mouse over the Remaining Capacity to view the reference VM specifications.
- The number of Host-Datastore combinations that have been excluded from the capacity analysis due to high total latency. You can click on **Excluded Host-Datastores**, to view the latency details of all the host and datastore combinations on the respective cluster or datacenter.

The screenshot displays the HPE Cloud Optimizer Capacity Modeler interface. It includes a navigation panel on the left, a main header with tabs for Overview, Performance, Capacity, and Groups, and a central Capacity Modeler form. The form contains input fields for VMs, vCPUs, Physical RAM, Disk, and Memory, along with start and end dates. Below the form is a table showing the results of the capacity model, including columns for ID, Action Name, Entity Name, Start Date, End Date, and a Result column. The table shows a single entry for 'Add new VMs' with a 'Failed - CPU, Memory' result. Below the table are three colored boxes (CPU, Memory, Storage) indicating status like 'Shorfall seen after VM added from today'. At the bottom, there are four line graphs showing resource usage over time for CPU, Memory, and Storage.

Legend	Description
1. Select Feature	Provides the HPE Cloud Optimizer features.
2. Browse using the Navigation panel	For VMware domain, select the Datacenter or Cluster. For KVM domain, select the Business Group of KVM hosts.
3. Add Action	Provide the attributes of the resources for a Datacenter, Cluster, or Business Group to increase the workload of your environment.
4. Run Scenario	Check the availability of the resources for a specified date. If the resources are available you can allocate the resources else re-plan the allocation.

Legend	Description
5. Show Forecast	Determine the value of CPU, Memory, or Storage utilization after 30, 60, and 90 days of adding the resources.
6. View the Forecast Graphs	Determine the maximum shortfall for CPU, Memory, or Storage utilization.

You can provide data of the resources using the Modeler Input and analyze the impact of resource allocation using the Modeler Output. For more information, see [Using the Modeler Input](#) and [Using the Modeler Output](#).

Modeler Input

To use the Modeler Input functionality, you must have the number and attributes of the entities to be added or removed from the environment. The tabular data on the right side of the page displays the total count of the hosts, VMs, remaining capacity, Memory and CPU capacity and datastores in the particular Datacenter, Cluster, or Business Group. With the help of this information it is easier to allocate the resources in your environment.

Using the Modeler Input you can:

- [Add New VMs](#)
- [Add Cloned VMs](#)
- [Delete Host](#)

Once you add the required details of the VM for a particular datacenter or cluster you can add the actions to the table and validate the details. After validating the details, you can analyze the impact of adding or removing the VM using the Modeler Output functionality.

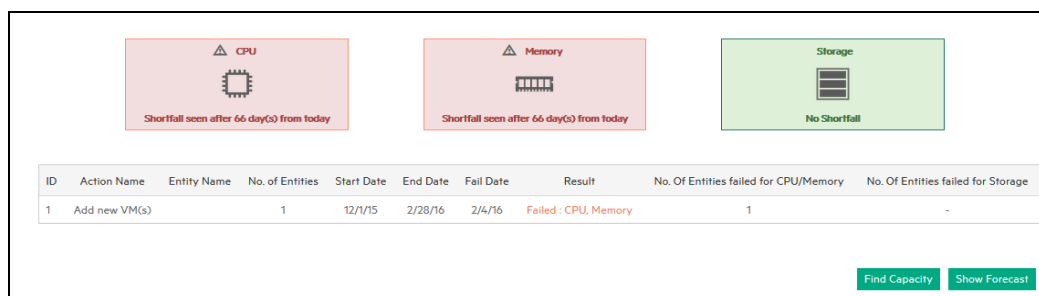
Modeler Output

After you input all the required details, HPE Cloud Optimizer provides you the forecast information for CPU, Memory, and Storage utilization of the environment. You can check if the resources allocated for a specific date are available or not. Based on the result, you can either allocate or re-allocate your resources for a particular datacenter or cluster.

You can use the tabular and graphical information to analyze how the environment performs after increasing the workload.

The boxes display the predicted CPU, Memory, and Storage shortfall based on the executed scenario. You can see the day at which the shortfall is first observed for the resources.

The table under the boxes displays the result of the individual scenarios. If the shortfall of resources is observed on any day during the 90 day interval starting from the current date, the scenario is considered as failed.



The graphs represent the CPU, Memory, and Storage utilization after increasing the workload. You can use the graphs to assess the trend of the utilization. If you see a steep rise or drop, it indicates that the performance of your environment is affected by the additional workload. The graphs give the following details:

- Days to Capacity - The number of days in which the CPU, Memory, or Storage of the environment reaches capacity or maximum value.
- Maximum Short Fall - The maximum difference between capacity and usage of CPU, Memory, or Storage in the environment over time.

Actions

Based on the information provided by the tables and graphs, you can determine whether increasing the workload with the current allocation is good for your environment or not. If the tables and graphs show that it impacts your environment, you can try a different allocation and see the results again.

Scenarios

Different capacity modeling scenarios are supported in HPE Cloud Optimizer. You can choose the required scenario, provide the entity details, and view the output.

Add New VMs

To increase your workload you may want to add more VMs to your environment. The VM attributes can vary based on the requirement. You can use the Modeler to provide the VM details and view the output. You can add multiple scenarios by varying the attributes of the requirement and then perform the analysis.

Add Cloned VMs

You may want to add VMs to your environment, which have same or similar attributes of existing VMs in the environment. For this, you can use the Add Cloned VMs scenario. You can also specify any attributes that need to be different from the existing VMs.

Delete Host

You may want to delete a particular host to distribute or decrease the workload of your environment. If the load on a host is low, you can re-allocate the VMs and analyze the


impact on the environment. During the maintenance period, you can add, move or remove the host for a particular time period to check the impact on the environment. Using the Modeler functionality, you can perform analysis by excluding the host for a certain period and checking the capacity of the datacenter.

Adding New/Cloned VMs or Deleting a Host

To add a new/cloned VM, or to delete a host, follow these steps:

1. Select the Datacenter, Cluster, or Business Group from the **Environment** tree on the left pane.
2. Select **Add new VM(s)/ Add cloned VM(s)/ Delete Host** scenario from the **Select Action** drop-down list.
3. Based on the scenario selected, add the following values:
For Add new VM(s): **No. of VMs, Start Date, No. of vCPUs, End Date, Physical RAM(GB), CPU Demand(GHz), Disk(GB), and Memory Used(GB)** fields.
For Add Cloned VM(s): **No. of VMs, Start Date, End Date, No. of vCPUs, and Physical RAM(GB)**
For Delete Host: **Start Date and End Date**
4. Click **Add Action** after verifying the data provided.

Note: Click **Set Defaults** to reset the default values.

The VM information appears in the **Scenario Workflow** table. From the Scenario Workflow table, you can click  , to delete the entry in the table.

5. Click **Run Scenario**. The CPU, Memory, and Storage utilization of the VM is displayed.

A table appears displaying the availability of the resources for the given date. If the resources can be utilized it passes the test else it fails the test.

6. Click **Show Forecast**. The forecast graphs for CPU, Memory, and Storage utilization for the particular datacenter is displayed.

A table appears on the right side of the graph displaying **Days to Capacity** and **Maximum Short Fall**.

Find Capacity

HPE Cloud Optimizer provides you a simple and intuitive interface named **Find Capacity** that helps you find compute capacity in your virtualized VMware environments. HPE Cloud Optimizer's capacity modeling capabilities combined with the Find Capacity functionality helps you to achieve end to end capacity planning and provisioning.

The Find Capacity feature allows you to find hosts to get additional CPU and Memory capacity when there is a shortfall of resources. If HPE Cloud Optimizer is integrated with HPE OneView, you can find additional capacity in both un-configured Server Hardware and Hypervisor Hosts. The Hypervisor Hosts are listed only if the shortfall cannot be offset using the Server Hardware capacity. You can launch HPE OneView directly from the Find Capacity interface and configure the Server Hardwares.

Navigation

You can access the Find Capacity feature in the following way:

From the Capacity Modeler

To access **Find Capacity** from the **Capacity Modeler**, perform the following steps:

1. From the HPE Cloud Optimizer home page, select **Capacity > Modeler**.
2. Run a modeling scenario by selecting the required action. For more information, see [Modeler](#).
3. If there is a shortfall of CPU or memory during the capacity modeling scenario, the **Find Capacity** button appears.
4. Click the **Find Capacity** button to access the Find Capacity interface.

Using the Find Capacity Feature

1. **Setting the Inputs**
 - a. In the **Shortfall** box, enter the amount of CPU and Memory resources that you want to borrow from other hosts in the environment. You can click **Details**, to view the details of the failed scenario.
 - b. In the **Duration** box, select the time interval for which you want to borrow CPU and Memory resources from other hosts in the environment.
 - c. In the **Container** box, select the container in which you want to find Hosts. The container can be **DataCenter** or **Cluster**.
 - d. Click **View/Select Containers**. The **View/Select Containers** dialog box appears.

The **View/Select Containers** dialog box contains the list of available containers in the environment from which you can borrow resources. It also lists containers that cannot be selected as their Days To Capacity (DTC) values are less than the end date selected. The DTC is the number of days in which the entity reaches the maximum capacity. You can also access the Capacity Usage Trends, Workbench, and Forecast for the individual containers.
 - e. Select the **Include Un-configured Server Hardware** check-box if you want to find capacity in the un-configured Server Hardware.

Note: The **Include Un-configured Server Hardware** check-box appears only when HPE Cloud Optimizer is integrated with HPE OneView.

2. Click **Find Host**. A list of un-configured Server Hardware and Hosts appears.

3. **Analyzing the Results**

a. The lists of Hosts appears with the following information:

Field	Description
Host Name	Name of the Host that can be borrowed. Click the Host Name to view the Host Summary information.
CPU Capacity (GHz)	The CPU capacity of the Host.
Mem Capacity (GB)	The memory capacity of the Host.
Cluster Name	The name of the Cluster to which the Host belongs to. Click the Cluster Name to view the Cluster Summary information.
vCenter Name/Datacenter Name	The vCenter and Datacenter name to which the Host belongs to.

b. If you have selected the **Include Un-configured Server Hardware** check-box, you can see the list of un-configured Server Hardware in the environment with the following information:

Field	Description
Server Hardware	The name of the un-configured Server Hardware. Click the Server Hardware name to launch HPE OneView.
CPU Capacity (GHz)	The CPU capacity of the Server Hardware.
Mem Capacity (GB)	The memory capacity of the Server Hardware.
Enclosure	The name of the Enclosure on which the Server Hardware is available. Click the Enclosure name to launch HPE OneView.

Use Case

Bill is a virtualization Subject Matter Expert (SME) in an organization. Bill needs to add 50 VMs to his vCenter. Before taking the action, Bill needs to ensure that the virtual infrastructure is capable of handling the additional VMs.

This use case demonstrates how Bill uses the Capacity Modeler and the Find Capacity feature to accomplish his objectives.

The tasks that Bill must perform as an Administrator are as follows:

1. Log on to the HPE Cloud Optimizer.
2. Access the **Capacity Modeler** from **Capacity > Modeler**.
3. In the Capacity Modeler specify the details of adding 50 VMs to a Datacenter or Cluster. For more information on using the Capacity Modeler, see [Modeler](#).
4. If the Capacity Modeler displays shortage of CPU or Memory resources, the **Find Capacity** button appears.
5. Click **Find Capacity**. The **Find Capacity** interface appears with the list of un-configured Server Hardware or hypervisor hosts from which you can borrow capacity.

Based on this information, Bill can perform the following actions:

- Configure the listed un-configured Server Hardware in HPE OneView and offset the shortfall of resources, or
- Remove the listed Host from the cluster to which it belongs and add it to the datacenter or cluster that is facing the shortfall of resources.

Configuring Capacity Head Room Settings

Capacity Head Room is defined as the buffer or reserved capacity that is set aside to avoid the capacity challenges in the environment.

Monitored Capacity = Total Capacity - Capacity Head Room

Configuring Capacity Head Room will affect the following features:

- Capacity Overview
- Capacity Modeler
- Alerts

You can configure the head room settings by any *one* of the following methods:

- [Updating the Head room levels in pvcd . PO namespace](#)
- [Using the command line](#)

Configure the Head Room Values in pvcd . PO namespace

Follow the steps:

1. Log on to the machine as root.
2. Run the command:
/opt/OV/bin/ovconfchg -edit
3. Open the `pvcd.PO` namespace and update the following values:
`RIGHTSIZE_CPU_HEADROOM=<Value>`
`RIGHTSIZE_MEM_HEADROOM=<Value>`
`RIGHTSIZE_DISK_HEADROOM=<Value>`
In this instance,
the *Value* is the head room value for CPU, memory, and disk.
By default, the head room value for CPU and memory is 20, disk is 10.

Using the Command Line

Follow the steps:

1. Log on to the machine as root.
2. Run the following command to update the head room values for CPU and Memory:
/opt/OV/bin/ovconfchg -ns pvcd.PO -set RIGHTSIZE_CPU_HEADROOM=<Value>
/opt/OV/bin/ovconfchg -ns pvcd.PO -set RIGHTSIZE_MEM_HEADROOM=<Value>
/opt/OV/bin/ovconfchg -ns pvcd.PO -set RIGHTSIZE_DISK_HEADROOM=<Value>
In this instance, the *Value* is the head room value for CPU, memory, and disk. By default, the head room value for CPU and memory is 20, disk is 10.

Chapter 9: Troubleshooting HPE Cloud Optimizer

The following section details how to troubleshoot HPE Cloud Optimizer:

Unable to scroll to bottom of browser window

Symptom	When accessing HPE Cloud Optimizer user interface, you are unable to scroll to the bottom of the browser window.
Cause	This happens when your screen resolution is not set to an appropriate value.
Resolution	Set your screen resolution to a higher value. The recommended screen resolution for viewing the HPE Cloud Optimizer interface is 1280x768 and the browser must be in the full screen mode.

Unable to log on to the HPE Cloud Optimizer interface after logging out

Symptom	When you log out and try to log on to the HPE Cloud Optimizer user interface again, the page refreshes to display the Login page again.
Resolution	Provide your credentials and try to log on again. The HPE Cloud Optimizer user interface page appears.

Unable to View Graph in Certain Formats (XLS/TSV)

Symptom	After drawing a performance graph, you are not able to view it in formats like .xls and .tsv.
Cause	Option to download files in the Microsoft Excel and .tsv format may not be enabled in the browser security settings.
Resolution	<p>To check the browser settings, perform the following steps:</p> <ol style="list-style-type: none">1. Select Internet Options from the Tools menu of IE browser. The Internet Options window opens.2. Click the Security tab. Check if the security level in your browser settings is set to High.3. Make sure the Internet icon is selected and click Custom Level. The Security Settings window opens.

	<ol style="list-style-type: none"> 4. Scroll down to the Downloads section and select the Enable option under File Download. 5. Click OK and close the Internet Options window. 6. Open the Control Panel and double-click Folder Options. 7. Click the File Types tab to display the registered file types. 8. Select XLS file format and TSV file format from the Registered file types list, and then click Advanced. The Edit File Type window opens. 9. Click to clear the Confirm Open After Download check box. 10. Click OK.
--	---

Unable to connect to VMware vCenter Server

Symptom	After installing HPE Cloud Optimizer, you are not able to access HPE Cloud Optimizer from VMware vCenter Server.
Cause	There may be a firewall between HPE Cloud Optimizer server and VMware vCenter Server.
Resolution	Configure the firewall to allow communication between HPE Cloud Optimizer server and the VMware vCenter Server. The ports that must be opened are 80 and 443.

Unable to access HPE Cloud Optimizer in BSM when it is accessed using HTTPS

Symptom	When HPE Cloud Optimizer is accessed in MyBSM using HTTPS, the browser displays a security warning: ' Do you want to view only the webpage content that was delivered securely? '. If you click Yes , the webpage displays the message: ' Navigation to the webpage was canceled. '
Cause	The security settings of the browser are set such that the mixed content warnings are enabled.
Resolution	<p>To disable the mixed content warning:</p> <ol style="list-style-type: none"> 1. In the Internet Explorer, go to Tools > Internet Options > Security. 2. Select the Internet icon. 3. Under Security levels for this zone, click Custom Level. 4. In the Security Settings window, go to Miscellaneous > Display mixed content.

	<p>5. Click the Disable radio button.</p> <p>6. Click OK.</p> <p>The mixed content warnings are disabled.</p>
--	---

Collection is not working in KVM/Xen

Symptom	Data is not collected for one or more datasources.
Cause	<ul style="list-style-type: none"> ssh is not automated The trace logs in pvtrace.0.txt are not logged properly. There are 2 opsagt processes running on the HPE Cloud Optimizer virtual appliance and one opsagt process is the child of another.
Resolution	<p>To check whether collection is not working:</p> <ul style="list-style-type: none"> Check whether ssh is automated. To check whether ssh is automated: <ul style="list-style-type: none"> a. Run the command: <pre>ssh <username>@<KVM/Xen hostname></pre> where, <i><username></i> is the username of the KVM/Xen user and <i><KVM/Xen hostname></i> is the hostname of the KVM/Xen host. b. If you are prompted for a password, then the collection is not working. Run the command: For KVM: <code>virsh -c qemu+ssh://root@<hostname>/system</code> For Xen: <code>virsh -c xen+ssh://root@<hostname>/</code> Where, <i><hostname></i> is the host name of the KVM/Xen server. If you are prompted for a password, then the collection is not working. Check for trace logs in pvtrace.0.txt located at /var/opt/OV/log/. <p>Solution 1</p> <p>To restart the collection:</p> <ul style="list-style-type: none"> Click Settings > Collection & Analysis. In the Status of Datasources table, from Actions, click Restart. <p>Solution 2</p> <p>If there are 2 pvcd processes running on the HPE Cloud Optimizer virtual appliance and one pvcd process is the child of another:</p>

	<ol style="list-style-type: none"> 1. Stop all the pvcd processes that are running, using the following command: <code>ovc -stop pvcd</code> 2. It is recommended that you update the libvirt library to the latest version on the HPE Cloud Optimizer machine. 3. Restart collection using the following command: <code>ovc -start pvcore</code>
--	---

Unable to launch Real Time Guest OS Drilldown, port is not available

Symptom	When accessing HPE Cloud Optimizer user interface, you are unable to connect and open the Guest OS Drill down page as the default port is not free.
Cause	This happens when the port is already in use.
Resolution	<p>To configure the port number:</p> <ol style="list-style-type: none"> 1. Log on to the node as administrator. 2. Go to the directory: <ul style="list-style-type: none"> • On Linux or Linux Debian node - <code>/opt/ OV/ hpcs</code> 3. Open the hpcs.conf file and edit the default port number and collection interval in the hpcs.runtime namespace. 4. Modify the default port number as port=<Value>. By default, nodes deployed with Real Time Guest OS Drilldown component use the port 381 for communication. In this instance, <Value> is the port number that is used by Real Time Guest OS Drilldown component. 5. Restart Real Time Guest OS Drilldown component.

Trace messages are not getting logged for Capacity Modeler and RTAD during runtime

Symptom	The trace messages for Capacity Modeler and RTAD is available only on the interface when run from the command prompt.
Resolution	<p>To enable runtime tracing for Capacity Modeler and RTAD, follow these steps:</p> <ol style="list-style-type: none"> 1. Create a trace config file (<code>/tmp/pvcd.tcf</code>) with the following content:

	<p>TCF Version 3.2</p> <p>APP: "pvcd"</p> <p>SINK: File "/tmp/pvcd.tr" "force=0;maxfiles=10;maxsize=10;"</p> <p>TRACE: "oafmrwrk.data" "Trace" Info Warn Error Developer Verbose</p> <p>TRACE: "whatif" "Trace" Info Warn Error Developer Verbose</p> <p>TRACE: "thresholdtrigger" "Trace" Info Warn Error Developer Verbose</p> <ol style="list-style-type: none"> 2. Run the command <code>/opt/OV/support/ovtrccfg -cf /tmp/pvcd.tcf</code> 3. Access the Capacity Modeler from the interface or wait for RTAD to complete a collection. 4. A trace file <code>/tmp/pvcd.tr_00000.trc</code> is created with the trace contents. 5. Run the command <code>/opt/OV/support/ovtrccfg -off</code> to switch off tracing.
--	--

Metrics are missing when a graph is viewed as a table

Symptom	<p>In the workbench window:</p> <ol style="list-style-type: none"> 1. Go to the Graphs tab, and generate a graph using multiple metrics. 2. From the Options menu in any graph, select the View as Table option. <p>The table view for the graph does not display all the metrics selected to generate the graph.</p>
Cause	<p>The metrics are available in the table, but they are not visible because the horizontal scroll bar for the table is hidden.</p>
Resolution	<p>To view all the metrics:</p> <ol style="list-style-type: none"> 1. Scroll down the table using the vertical scroll bar. 2. A second horizontal scroll bar - in addition to the scroll bar already visible - appears. <p>Scroll across the table using this scroll bar.</p>

Data does not display for graphs in Capacity Usage Trends

Symptom	<ol style="list-style-type: none">1. CPU usage and memory usage values for Cluster displays as 0 in the Resource Information table on the Capacity Usage Trends page.2. Storage Allocation graph does not show any data in the Overview page.3. The Utilization by File Types graph does not display any data in the Summary for Data Stores page.
Cause	HPE Cloud Optimizer server time and vCenter time may not be in sync.
Resolution	Check if the HPE Cloud Optimizer server time and datasource time are in sync. HPE Cloud Optimizer collects time bound data from the vCenter. So time must be synchronized across the HPE Cloud Optimizer virtual appliance and the vCenter appliance or server.

Not able to access HPE Cloud Optimizer user interface when PKI authentication is enabled

Symptom	When trying to access HPE Cloud Optimizer while PKI authentication is enabled, browser displays the message "Page cannot be displayed"
Cause	The CAC certificate might have expired or is revoked.
Resolution	Check if the certificate is expired or revoked.

Tooltip not appearing adjacent to the graph

Symptom	When accessing HPE Cloud Optimizer in Firefox, the tooltips for the graphs are not appearing close to the cursor.
Cause	The offset value for HTML elements in different browsers are different.
Resolution	Access HPE Cloud Optimizer using Google Chrome.

Graphs show less than 12 data points (based on 5 minute data collection) for an hour

Symptom	Based on 5 minute data collection, graphs must show 12 data points for an hour. But graphs are showing less than 12 data points.
---------	--

	<p>OR</p> <p>Data summarization slider in the Options pane of the Treemap does not display uniform data points for 5 minutes.</p> <p>OR</p> <p>Treemap shows data only for Hosts and VMs.</p>
Cause	<p>The reason can be one of the following:</p> <ul style="list-style-type: none"> • HPE Cloud Optimizer renders graphs based on summarized values, with a 5-minute interval. The data collector may not complete the collections on exact 5 minutes always. So 2 data points might get averaged out during summarization. <p>OR</p> <ul style="list-style-type: none"> • In large vSphere environments with high number of instances (more than 3000), data collection and logging in HPE Cloud Optimizer might get delayed. As a result, you may experience data gaps. Performance data collection for large number of instances is a resource and time intensive job for the vCenter server and HPE Cloud Optimizer. The huge volume of data, might result in the delay in data logging due to indexing and co-relation between entities.
Resolution	<p>If data gaps persist for long or you see data points are missing for more than 3 intervals, check the vCenter configuration, user permissions, and deployment environment of HPE Cloud Optimizer.</p>

PV Collection Daemon does not start after rebooting

Symptom	<p>PV Collection Daemon does not start after reboot. The <code>ovc -status</code> shows <i>pvcd</i> in aborted state, after rebooting.</p> <p>OR</p> <p>Vertica Database does not start after rebooting.</p>
Cause	<p>If the HPE Cloud Optimizer Virtual Appliance is shutdown abruptly, sometimes the database does not start. Due to this, <i>pvcd</i> exits on rebooting.</p>
Resolution	<p>Restart the <i>pv</i> process (<i>pvcd</i>, Tomcat, and Vertica) after rebooting the machine. Run the command, pv restart from the HPE Cloud Optimizer console.</p>

Tracing for Advanced Troubleshooting

Component	Description	Location
pvcd (PV Collection Daemon)	<p>pvcd daemon uses XPL tracing mechanism. A default pvcdXPL.itc (trace configuration) is available in /opt/OV/support directory.</p> <p>To start the tracing, run <code>/opt/OV/support/ovtrccfg -cf /opt/OV/newconfig/conf/xpl/trc/pvcdXPL.tcf</code> The XPL trace files are present at /tmp/pvcd_*.trc.</p> <p>To stop tracing, run <code>/opt/OV/support/ovtrccfg -app pvcd -off</code></p>	<p>If tracing is enabled, /tmp/ directory contains the trace files (use - ls /tmp/pvcd_*.trc)</p>
HPE Cloud Optimizer Web Server	To start tracing – pv trace on	/var/opt/OV/log
	To stop tracing – pv trace off	/var/opt/OV/log

Data Collection Status Messages

In HPE Cloud Optimizer, when data collection does not happen properly, the system displays some error messages. The following table lists the error messages, and scenarios when these messages can appear, and resolution, wherever applicable.

Error Message /Status Message	Scenario	Resolution
Connecting...	HPE Cloud Optimizer is establishing connection with the datasource.	No resolution required.
Connected	Collection is completed successfully and waiting for the next collection to	No resolution required.

Error Message /Status Message	Scenario	Resolution
	start.	
Data Collection in Progress...	Collecting data and storing it in the database.	No resolution required.
Total number of instances exceed the maximum limit of the free license instance count (<maximum count>). Install a permanent license with sufficient capacity or enable the evaluation license from the Settings page.	Instance count exceeds the specified limit of the free license.	<p>Perform one of the following actions:</p> <ul style="list-style-type: none"> • Wait for the next data collection to complete. • Remove a datasource to adjust the instance count. For more information on removing datasources, see Removing a Datasource. • Install a permanent license with sufficient capacity or enable the evaluation license from the Settings page.
Total number of instances exceed the maximum limit of the license instance count (<maximum count>). Use a datasource or datasource credentials with fewer instances associated.	Instance count exceeds the specified limit of the free license.	<p>Perform one of the following actions:</p> <ul style="list-style-type: none"> • Use a datasource or datasource credentials with fewer instances associated. • Remove a datasource to adjust the instance count. For more information on removing datasources, see Removing a

Error Message /Status Message	Scenario	Resolution
		Datasource.
Total number of instances exceeds the maximum limit of the license instance count (<maximum count = 25>). Wait for the next data collection to complete or restart using the 'Restart Collection' button.	Instance count exceeds the specified limit of the free license.	Perform one of the following actions: <ul style="list-style-type: none"> • Wait for the next data collection complete. • Go to Settings > Collection & Analysis. In Status of Datasources table, from Actions, click Restart.
Total number of Hosts and VMs exceeds the maximum count 25. Use vCenter credentials with fewer Hosts and VMs associated.	Instance count exceeds the specified limit.	Perform one of the following actions: <ul style="list-style-type: none"> • Remove a datasource to adjust the instance count. For more information on removing datasources, see Removing a Datasource. • Limit the number of instances associated with a datasource.
Data Collection failed. Check if Libvirt services are running on the HPE Cloud Optimizer server.	The Libvirt services are not running on the HPE Cloud Optimizer server.	The Libvirt services must be running on the HPE Cloud Optimizer server.
Invalid username or password	Credentials provided are wrong.	Provide the correct credentials.
Connection Timeout	Connection request to the	Wait for the next data collection.

Error Message /Status Message	Scenario	Resolution
	datasource has timed-out.	
Unable to connect to vCenter server	<p>This can be due to any of the following scenarios:</p> <ul style="list-style-type: none"> • Unable to connect to datasource • Invalid datasource name • System is not on the network. 	<p>Actions:</p> <ul style="list-style-type: none"> • Check if the datasource is running. • Provide the correct datasource name. • Ensure that the datasource is accessible from HPE Cloud Optimizer.
Unable to collect performance data from the vCenter	No performance data is coming from the datasource.	Make sure that the datasource services are up and running and historical data collection is happening.
Unable to connect to the Database	Database connection fails.	Try restarting HPE Cloud Optimizer.
Unable to start collection: Error occurred while connecting to the Database	Database is not running or HPE Cloud Optimizer is not able to connect to the database.	Try restarting HPE Cloud Optimizer.
Unable to start collection: MetricList is empty	Metric list in the integration XML is empty.	Contact Support.
Unable to start collection: Error occurred while reading Database configuration	Error occurred	Contact Support.

Error Message /Status Message	Scenario	Resolution
	while reading the database configuration file.	
Unable to start collection: Model generation failed	Database model generation fails.	Contact Support.
Not Connected	Collection threads fail to get created or are yet to be created.	Contact Support.
Connection Error	Generic message shown when there is an unknown exception.	Check <i>logs/trace</i> .

Unable to install the collector service on Windows 2012 R2

Symptom	<p>After downloading the Hyper-V proxy collector (vPVWinVirtCollector.zip) on Windows 2012 R2, you get the following error while installing the Install.bat collector.</p> <p>The system cannot find the path specified.</p> <p>"Installing the vPV Collection Service"</p> <p>Microsoft (R) .NET Framework Installation utility</p> <p>Version 4.0.30319.33440</p> <p>Copyright (C) Microsoft Corporation. All rights reserved.</p> <p>Exception occurred while initializing the installation:</p> <p>System.IO.FileLoadException: Could not load file or assembly 'file:///C:\New folder\vPVWinVirtCollector.exe'</p>
---------	---

	<p>or one of its dependencies. Operation is not supported. (Exception from HRESULT: 0x80131515). "Starting the vPV Collection Service" The service name is invalid. More help is available by typing NET HELPMSG 2185</p>
Cause	While installing Install.bat collector, you get an error due to the group security policy.
Resolution	<p>To remove the group security policy, follow these steps:</p> <ol style="list-style-type: none"> 1. Go to the directory where you have extracted the vPVWinVirtCollector.zip package. 2. Right-click the vPVWinVirtCollector.exe, system.management.automation.dll, and JsonSharp.dll respectively and go to Properties. 3. In the General tab, go to Security field and click Unblock. 4. Run the Install.bat as a domain user with administrative privileges to start the collection.

HPE Cloud Optimizer loses connection to vCenter

Symptom	Data collection stops for vCenter.
Cause	HPE Cloud Optimizer loses connection to vCenter target.
Resolution	<p>Run the following command to restart the collection:</p> <pre>ovc -restart pvcd</pre>

Hyper-V Collector fails to upload data to HPE Cloud Optimizer

Symptom	Hyper-V Collector fails to upload data to HPE Cloud Optimizer.
Resolution	<p>Set the recovery option for HPE Cloud Optimizer collector service. To set the Recovery option:</p> <ol style="list-style-type: none"> 1. Open the Start menu and type services.msc in the Search box. 2. Press Enter. The Services window appears. 3. Right click HPE Cloud Optimizer Collector Service and click

	<p>Properties. The HPE Cloud Optimizer Collector Service Properties dialog box appears.</p> <p>4. Select the Recovery tab.</p> <p>5. In the First Failure list, select Restart the Service.</p> <p>6. Click Apply.</p>
--	---

Data collection takes more than five minutes in some vSphere environments.

Symptom	Data collection takes more than five minutes in some vSphere environments.
Resolution	<p>Check the overall performance of the datasource from the Datasource Health page. If the Status of the datasource is Slow Responsive Target or the Query Time Taken is consistently higher than 280 seconds, it indicates poor performance of the datasource.</p> <p>Identify the possible reasons for poor performance and take necessary corrective actions to improve the responsiveness. Some of the possible reasons can be:</p> <ul style="list-style-type: none"> • High resource utilization in the vCenter due to multiple client connection sessions • Undersized vCenter • Over-utilization of the vCenter • ESXi Host where vPV Virtual Appliance is deployed being saturated for CPU, memory, or storage • Undersized vPV Virtual Appliance <p>If the response time for the vCenter could not be improved by corrective actions, consider changing the data collection interval to 15 minutes.</p>

Unable to open HPE Cloud Optimizer reports from HPE Operations Manager / (OMi) Events Browser

Symptom	When HPE Cloud Optimizer is integrated with OMi, you can view the HPE Cloud Optimizer events in the OMi Event Browser. However, you are not be able to launch HPE Cloud Optimizer reports from the OMi Event Browser.
Resolution	To view the reports, copy the report URL and paste it in a browser.

Data collection for Physical Server in HPE Cloud Optimizer fails in a large environment (>4000 instances)

Symptom	Data collection for Physical Server in HPE Cloud Optimizer fails in a large environment (>4000 instances).
Resolution	<p>Perform the following steps to resolve the issue:</p> <ol style="list-style-type: none"> Run the following command to stop ovc. <code>ovc -stop</code> Ensure that all the ovc processes are stopped. Run the following commands: <code>/opt/OV/bin/ovconfchg -ns bbc.http.ext.bbc.cb.ovbbccb -set MAX_SOCKETSET_THREADS 30</code> <code>/opt/OV/bin/ovconfchg -ns bbc.http.ext.pvc.pvcd -set MAX_SOCKETSET_THREADS 30</code> Run the following command to start ovc. <code>ovc -start</code>

In the Capacity Overview for VMware, the storage graphs for a few hosts appear as blank.

Symptom	In the Capacity Overview for VMware, the storage graphs for a few hosts appear as blank.
Resolution	<p>This behavior can be observed due to some configuration or hardware changes for the host. To resolve the issue, run the following command at the command prompt:</p> <code>pv restart</code>

Unable to upgrade HPE Cloud Optimizer from 2.20 to 3.00 using CHF4 from VA

Symptom	<p>[ERROR] Failed with exit code 20480</p> <p>Update status: Error while running installation tests</p> <p>[ERROR] Failure: updatecli exiting abnormally</p>
Resolution	<p>The behavior can be observed due to unavailability of space in the / partition.</p> <p>Clean up the / partition to resolve the issue.</p>

Unable to set the threshold value for Datastore Space Utilization alert

Symptom	No parameters found in namespace [DatastoreSpaceUtilMonitor] in policy vPV-CustomAlertSensitivityDefinition to set the threshold value for datastore space utilization monitoring.
Cause	RIGHTSIZE_DISK_HEADROOM variable is not exposed in the vPV-CustomAlertSensitivityDefinition policy.
Resolution	<p>Add the following to the vPV-CustomAlertSensitivityDefinition policy and deploy it to the HPE Cloud Optimizer system:</p> <pre>[pvcd.P0] RIGHTSIZE_DISK_HEADROOM=<></pre> <p>Example:</p> <p>Threshold value for datastore space utilization = 100 - RIGHTSIZE_DISK_HEADROOM. By default, the parameter RIGHTSIZE_DISK_HEADROOM is set to 10. So, the default threshold is 90. To change the threshold value to 75,</p> <ol style="list-style-type: none"> 1. Edit the vPV-CustomAlertSensitivityDefinition policy. 2. Update RIGHTSIZE_DISK_HEADROOM as follows: <pre>[pvcd.P0] RIGHTSIZE_DISK_HEADROOM=25</pre>

Frequently Asked Questions

All Datasources

- Why do some graphs show a message "No data found for the duration with points every 5 minutes...." when I click the Show Performance context-menu in the Treemap?

For some metrics, the data does not get collected during the initial collection but is collected later. For the graphs using these metrics, you may not see data after the initial collection. But data will be available after a few collections. These metrics and graphs are as listed in the following table:

Class	Graph	Metric
Host	Disk Usage	disk_read_average_kiloBytesPerSecond
		disk_write_average_kiloBytesPerSecond
	Memory Savings	mem_sharedcommon_average_kiloBytes
		mem_shared_average_kiloBytes
Guest	Disk Usage	DiskPhysIOByteRate
		DiskPhysReadByteRate
		DiskPhysWriteByteRate

- Why does data collection always show as In progress and HPE Cloud Optimizer utilize too much CPU to run?

By default, HPE Cloud Optimizer collects data from the datasources every 5 minutes. If a datasource response is slow or HPE Cloud Optimizer is configured to monitor too many instances, the data collection may take more than 5 minutes.

- Why is data not available for some metrics?

For the following metrics there is no data collected from vCenter 4.1:

- **Guest:**

- Cpu_system_summation_milliseconds
- CPUUsedTime
- CPUWaitTime
- NetInPacket
- NetOutPacket

- **Host:**

- Cpu_used_summation_milliseconds
- Cpu_idle_summation_milliseconds
- NetInPacket
- NetOutPacket

Note: Data store metrics are collected only from vCenter 5.0 and above versions.

- Why am I not able to see folders in **HPE Cloud Optimizer**?

Folders are not available in the Treemap and Configuration Items tree, but the

corresponding VMs are added to the parent in the hierarchy.

- Some resources display the name of their immediate parents while other resources don't.

If multiple datasources of the same type contains instances with same name, then the HPE Cloud Optimizer Treemap displays the immediate parent name of the instances alongside the instance name. The instance name and parent name will be written as: *<instance_name (parent_name)>*.

For example:

Scenario 1

In the vCenter domain, if a VM (VM1) belongs to two different vCenters, then, in the Treemap, the host name (Host1) of the VM is displayed alongside the VM name. The VM and host names are displayed as *VM1 (Host1)*.

Scenario 2

In the vCenter domain, if a Host (Host1) belongs to two different vCenters, then, in the Treemap, the cluster name (C1) of the host is displayed alongside the host name. The host and cluster names are displayed as *Host1 (C1)*.

However, if Host1 does not belong to any cluster, then the datacenter name of the host is displayed: *Host1 (Datacenter)*.

Note: Similar nomenclature is used if the resources are grouped using the **Group By** option in the **Options** pane.

- Why is memory utilization always 100 % for KVM, Xen, and OpenStack?

The memory utilization is always 100 % for KVM, Xen, and OpenStack because currently the total memory capacity is used to denote the memory utilization in Libvirt. This is so because the current versions of Libvirt do not display the memory utilization.

- How do I check the status of the database in **HPE Cloud Optimizer**?

To check the status of the database:

- a. Run the following commands on the HPE Cloud Optimizer server:

```
su pv_vertica
/opt/vertrica/bin/adminTools
```

- b. The administration UI for the Vertica database opens.

You can check the status of the Vertica database in the administration UI.

- There is a mismatch between the instance count of a datasource displayed in the Treemap and the instance count of the same datasource displayed on the License and Collection & Analysis options in the Settings page.

When a datasource goes in an error state, the instances from that datasource are not considered while calculating the instance count displayed in the Settings page.

However, the instances from the erroneous datasource will still be considered while calculating the instance count displayed in the HPE Cloud Optimizer Treemap as long as the datasource is not deleted from the HPE Cloud Optimizer Settings page.

This is because the data from the Treemap is used to plot the graphs and reports in the workbench. Hence, it is necessary to have the information of all the datasources added in HPE Cloud Optimizer.

- The **HPE Cloud Optimizer** Treemap shows data for a datasource even if the datasource is in an error state.

The HPE Cloud Optimizer Treemap shows the data for a datasource in error state for one hour after the last collection time.

- The HPE Cloud Optimizer title bar hides the performance graphs' title bar if the performance graphs are dragged to the HPE Cloud Optimizer title bar. This prevents you from closing the performance graphs.

Refresh the browser to close the performance graphs.

- If I select the Highlight option from the Options pane, there is a mismatch between the VM count shown on the Treemap and the VM count shown on the Workbench.

The maximum valid Get URL length is 2000 characters. If the VM count in the highlighted resources is high, the overall Get URL length may exceed 2000 characters.

If the length of the VMs' Get URL exceeds 2000 characters, some of the URLs will be truncated. Hence, VMs with truncated URLs will not be visible on the Workbench.

- Why do I not see data for datastores in the Summary pages of the Optimization feature after starting HPE Cloud Optimizer?

If you do not see data for datastores, wait for one hour and then check the data again.

- Why does a VM name appear twice in the Treemap or Workbench?

When you recreate a VM with the same name, you can see two instances of the same VM for some time.

This is because while refreshing the Treemap or Workbench, the VM instances are considered for twice the data collection interval. For example, if the data collection interval is five minutes, all the VM instances that have data for an interval of ten minutes will be considered. As in this scenario, both the VM instances have data during that time interval, both the VM instances appear on the Treemap or Workbench. However, the earlier VM instance will not be seen after the data collection has happened twice.

- In the Forecast page, why does the resource usage trend for an entity (Host, Cluster, or Datacenter) show an increasing trend, when the resource usage for individual VMs is flat?

Such a trend is observed when the number of VMs for an entity is increasing over a period of time.

- What are the log files available with HPE Cloud Optimizer?

The following table lists the log files available with HPE Cloud Optimizer:

Component	Description	Location
pvcd (PV Collection Daemon) and pvconfig	pvcd and pvconfig tools log messages in System.txt.	/var/opt/OV/log/System.txt
Collectors	vCenter and OpenStack collector status log file /var/opt/OV/log/status.virtserver	/var/opt/OV/log/status.virtserver
	KVM collector status log file	/var/opt/OV/log/status.kvm
	XEN collector status log file	/var/opt/OV/log/tmp/status.xen
Database Transaction	Vertica catalog file to view database transactions	/var/opt/OV/databases/pv/catalog/pv/v_pv_node0001_catalog/vertica.log
HPE Cloud Optimizer Web Server	HPE Cloud Optimizer Web Server component logs messages in ovpm.0.log	/var/opt/OV/log/ovpm.0.log

VMware vCenter Server

- How can I set upper and lower threshold values for the Treemap colors so that I can see the color variation only for the range of my interest?

- a. Open the file
`OVINSTALLDIR/newconfig/OVPM/smepack/VCENTER/integration/VCENTER_GC_Integration.xml`.
- b. Locate the corresponding **CI_VIEW** tag (for example, VM (Compute)) and **METRIC** tag (for example, guestMemPercent) for which you want to set the thresholds.
- c. Add the additional tags- **COLOR_METRIC_MAX_VAL** for upper threshold value and **COLOR_METRIC_MIN_VAL** for lower threshold value.

For example,

```
<METRIC Name="MemEntlUtil" ColorCaption="UsagePercent"
SizeCaption="AvailGBs">
<COLOR_CLASS>GUEST</COLOR_CLASS>
<COLOR_METRIC>MemEntlUtil</COLOR_METRIC>
<COLOR_METRIC_MIN_VAL>50</COLOR_METRIC_MIN_VAL>
<COLOR_METRIC_MAX_VAL>80</COLOR_METRIC_MAX_VAL>
<SIZE_CLASS>GUEST_CONFIG</SIZE_CLASS>
<SIZE_METRIC>MemEntl/1024</SIZE_METRIC>
</METRIC>
```

According to the preceding example, any VM with **MemEntlUtil** ≤ 50 appears green and with **MemEntlUtil** ≥ 80 appears red. The color from green to red is distributed among VMs with value > 50 and < 80 for **MemEntlUtil**, giving more color options for that range.

Note: It is strongly recommended that you do not modify the XML. And if you modify, maintain the XML structure accurately. The changes are not retained during an upgrade. You may have to repeat the steps after upgrade.

- In VMware vCenter, first time data collection takes a long time

The initial collection time for VMware vCenter depends on:

- The number of instances from which data is collected.
- The response time of the VMware vCenter.

If the number of instances or the response time is high, then the initial collection will take long time to complete. The subsequent collections will take place quickly.

- After adding a new VMware vCenter Server to HPE Cloud Optimizer, why do I see a lot of items missing or no performance data items in the Treemap?

When you add a VMware vCenter Server, it takes some time for the entire data to get collected. Till the first data collection is complete, you may not see data in the Treemap, but the resource count in the Resource pane is updated. You must wait for

the first data collection to complete (last collection time gets updated in the **Collection & Analysis** option of the **Settings** page) to see the resources of the VMware vCenter Server in the Treemap.

- Why does not Datastore metric appear in HPE Cloud Optimizer?

This can happen when your VMware vCenter Server version is 4.x. Data store metrics are not available for version 4.x. Log on into VMware vCenter Server using the VI client and select the first element available in the tree in VI Client. The VMware vCenter Server details along with the version appear on the right pane.

- Why do I not see any micro charts for datastores?

You may not see micro charts for datastores if the VMware vCenter Server is just added. Storage metrics are collected every 30 minutes, so there is no data to display initially. You need to wait for two or more data collection cycles to complete.

- Which metrics are not collected if the Statistics Level of VMware vCenter is set below 2?

The following metrics are not collected if you set the Statistics level less than 2:

Hosts and Virtual Machines

- CPUCoStopTime
- CPUDemandUsed
- CPUIdleTime
- CPUIdleUtil
- CPUUsedTime
- CPUWaitTime
- DiskCommandAbortRate
- DiskDeviceLatency
- DiskKernelLatency
- DiskReadLatency
- DiskWriteLatency
- MemActive
- MemShared

- MemSharedCommon
- MemSwapped
- MemSwappedTarget
- MemBalloonTarget

Datastore

- DiskThroughputContention
 - DiskThroughputUsage
 - DiskReadRate
 - DiskWriteRate
- Why does not data store show up in Treemap?

Data store does not appear in the Treemap if performance data is not being collected for datastores at the VMware vCenter Server. Log on to the VMware vCenter Server using the VI client and check whether the **Performance** tab displays data for any of the datastores.

- Why do the default graphs for datastores not show any data?

Change the duration for the graph and check if data is shown in the graph.

- Why are resources available in the VI Client not visible in **HPE Cloud Optimizer**?

This may happen when historical data is not available. Make sure that historical data is available from the VI Client.

- I see data in Treemap even after removing all the VMware vCenter Servers.

The Treemap will get updated when the next data collection (for any VMware vCenter Server) added happens.

- Why do I not see Storage IO metrics in **HPE Cloud Optimizer**?

Storage IO metrics are not available for VMware vCenter Server version 4.x.

- What are the Logging levels required on a VMware vCenter Server?

The Logging levels 1 and 2 must be enabled on a VMware vCenter Server. Logging level 2 enables the VMware vCenter Server to collect data for the following Storage IO metrics:

- NumDiskReads
- NumDiskWrites
- When multiple VMware vCenter are deleted from the Status of Datasources table, the table still displays a few of the deleted VMware vCenters.

Close the Internet browser and open it again.

- Why do Reports and Capacity Usage Trends show mismatch in the allocation and utilization values of resources?

This is because the resource usage for hypervisor is also considered to calculate the utilization values. When the resource usage is closer to or has reached capacity, the usage values displayed exceed the allocation value as hypervisor usage is also included.

KVM/Xen

- In KVM and Xen domains, the memory and CPU metrics are not visible in the Metrics and Graphs tabs of the workbench.

The current version of Libvirt (0.10.2) does not support the memory and CPU metrics for Xen hosts. Hence, these metrics are not visible in the Metrics and Graphs tabs of the workbench. You can view these metrics using Domain-0 instance.

Also, for Xen hosts, the CPU utilization graph is not visible if the CPU utilization is less than one percent.

- Xen graphs do not show any data during the first five minutes.

HPE Cloud Optimizer does not collect the utilization and performance metrics for hosts and guests (VMs) during the first collection interval. Hence, the Xen graphs are not populated during the first collection interval.

- If any KVM/Xen host:
 1. Is disconnected by changing its IP address.
 2. And the same KVM/Xen host is reconnected again after some time.

All the guests in that datasource display CPU utilization as 100%

The guests will show the correct CPU utilization after 2 collection cycles.

Chapter 10: Real Time Guest OS Drilldown for Advance Troubleshooting

You can drill down in a Virtual Machine (VM) to troubleshoot the factors affecting the performance and health of the VM. This is the advance troubleshooting available with HPE Cloud Optimizer.

Real Time Guest OS drill down tool (HPE ComputeSensor) is a light-weight performance viewer tool that provides a quick overview of the workloads and applications that are available and running on the VM.

The tool provides you the following:

- Overview about the system resource utilization such as CPU, file systems, disks, and network utilization
- Overview about the general health of the system and reasons for any available bottlenecks
- List of processes and the resources that are being used
- System events generated in last 24 hours

You can determine if the system utilization is optimal or the health of the system by checking the graphs available. You can further drill down to detect the processes where the memory or CPU utilization is exceeding the limit by checking the processes that are consuming more CPU or memory.

To check the real time example, see [Using Guest OS Drill Down to Troubleshoot the Resource Utilization in your Environment](#).

Navigation

From the Treemap view, click **VM (Compute)** from the top pane. Select the VM and click **Guest OS Drill Down**, the Treemap refreshes to open the view to display resource utilization of the VMs.

If the Real Time Guest OS drill down feature is not available when you select a VM in Treemap, you can download the tool in the following ways:

- If you are an *administrator*, you can download it from the **Settings > Integration** page.
- or
- Follow the steps:

- a. From the **Treemap** view, click **VM (Compute)** from the top pane.
 - b. Select the VM and right-click. Select **Attempt Real Time Guest OS Drilldown**. The **Real Time Guest OS Drill down** page appears.
 - c. Click **Download**. Follow the instructions that appear on the page.
- If you are *not* an administrator, follow the instructions that appear on the page.
- For more information, see HPE Compute Sensor user guide.

Guest OS DrillDown User Interface

The following table lists the features on the Guest OS Drilldown page and the corresponding functions:

User Interface Details	Description
Graphs	Displays the graphs to check the resource utilization.
Tabs	Displays the details about the processes that are running, CPU usage, free disk space, swap memory, filesystem, and network usage running on VM.
Options	Contains the settings available on the Guest OS Drilldown page.

Graphs

The Real Time Guest OS drill down provides you the graphical overview of the health of the VM. The real time value is defined as the data collected in the interval of 1 -10 seconds. The default value is 2 seconds. The graphs get updated every 2 seconds as the collection interval is real time. You can change the interval from the available settings. For more information, see [Using the Options Pane](#).

You can use the graphs to check the utilization and analyze if the resources are adequately being used. You can derive the %utilization in the x interval. Here, x is the time in seconds.

The following graphs are available:

Graph Name	Description
CPU Utilization	Shows the real time CPU utilization.

Graph Name	Description
	<p>Here, the utilization reached around 45% at 15.48.20.</p> <p>You can derive the following:</p> <p>CPU usage is high at the given point so you would want to check the details from the processes tab and gather the reason for the spike in CPU utilization.</p>
Memory Utilization	Shows the real time Memory utilization.
File system	Shows the real time File system utilization.
Process	Shows all the processes running on the VM in the specific time period. The report shows the comparative analysis of the live processes that are running and the processes that are in queue when the CPU is busy.
Disk	Shows the real time disk utilization.
Network	Shows the real time Network utilization of the specific VM. You can get the incoming and outgoing bytes or packets for the selected interface.

Tabs

The tabs are as follows:

Process Cloud - Shows the significant processes running on the VM based on the workload of the VM. The process that are using more resources are depicted in the larger fonts. You can troubleshoot and check the reason of high utilization. As an administrator or capacity planner, you can find out the most significant processes running on the system.

You can click the **Options** to enable the process cloud settings.

The settings available are:

Field	Description
Process Cloud Settings	<p>Contains the option to enable the settings in the process cloud tab.</p> <ul style="list-style-type: none"> Size By Resource - Select the checkbox (CPU, memory or both) to view the processes where CPU, Memory or

	<p>both utilization are high.</p> <ul style="list-style-type: none">• Visible Processes - Set the option to view the n processes. Example, if you want to view only 10 processes, then move the bar from left to right.• Auto Refresh - By default, the checkbox is selected. The process cloud automatically gets refreshed. You can stop the refresh by de- selecting the checkbox. The process cloud will not change.
--	---

Here, the resource utilization is more for the process Oracle. You can also view the other processes that are using the resources but the significant usage is by Oracle. You can further drill down to check the processes usage and the process tab opens.

You can also view only n number of processes by using Process Cloud Settings option available.

Processes - Shows the parameters that are available with the processes running on the selected VM. Each process is provided with unique ID and name. You can view the CPU utilization and memory used in all the available processes. This helps in finding the processes where the resource utilization is high.

Parameters of the processes tab

Name	Description
Process ID	Unique identifier for the process.
Process Name	Name of the process.
Process Arguments	The full command line with which the process was initiated.
%Total CPU Utilization	The total CPU utilized.
Memory Used (MB)	The total memory utilized.
User	On Linux, the user is the name of a process or the login account (from /etc/passwd) of a process.

CPU - Shows the details of all the available CPU for the selected VM. The parameters provide additional information of the usage of real time CPU utilization of the specific

VM. You can view the total CPU utilization, user CPU utilization and system utilization. Each CPU is provided with unique ID.

Parameters of the CPU tab.

Name	Description
CPU ID	Unique identifier of the CPU.
CPU Utilization	The CPU usage of the selected VM.
Total Time (seconds)	The total time in user and system mode.
User Time (seconds)	The time that this CPU (or logical processor) was in user mode.
System Time (seconds)	The time that this CPU (or logical processor) was in system mode.
Interrupts	The number of device interrupts for this CPU during the interval.
CPU Clock (MHz)	The clock speed of the CPU in the current slot.

Disk - Shows the details of all the disks available on the specific VM. You can view the disk utilization such as disk physical read or write. Each disk is provided with unique ID and disk name.

Parameters of the Disk tab.

Name	Description
Disk ID	Unique identifier of the disk.
Disk Name	Name of the disk.
Physical Reads	The number of disk read is disk device during the interval.
Physical Writes	The number of disk writes per second during the interval.
Physical I/Os	The number of operations for this disk device during the interval.
Physical Read	The data transferred from this disk device during the interval.
Physical Write	The data transferred to this disk device during the interval

Name	Description
(KBs)	
Physical (KBs) Usage	The number of bytes transferred to or from this disk device during the interval.
Disk Time (Seconds)	The time, in seconds, that this disk device was busy transferring data during the interval.

File System - Shows all the File Systems available on the system. The parameters available will provide additional information of the usage of real time file system utilization of the specific VM. You can view the File System utilization such as space utilization, and space reserved.

Parameters of the File System tab

Name	Description
Device Number	Unique identifier of the device.
Device Name	Name of the device.
File System Directory	The directory defined for the file system.
%Space Utilization	The space utilized in percentage by the file system.
Space Used (GBs)	The space utilized by the file system.
Space Reserved (KBs)	The space reserved by the file system.

Network - Shows all the available network interfaces (physical or virtual) on the system. You can get the incoming and outgoing bytes or packets for the selected interface.

Parameters of the Network view.

Name	Description
Network ID	Unique identifier for the network.
Network Interface Name	Name of the interface.
Network In	The total number of bytes received on the network interface.
Network Out	The total number of bytes transmitted on the network interface

Name	Description
Input Packets	The total number of packets received on the network interface.
Output Packets	The total number of packets transmitted on the network interface.
%Network Utilization	The throughput of the network, in %.
Errors	The total number errors during transmission or while receiving the bytes or packets.

Application - Shows all the processes running on the system that are grouped or aggregated into applications. The processes listed are active and alive processes. You can view the CPU and memory utilization of each application.

Parameters of the Application Tab.

Name	Description
Application Number	Unique identifier for the application.
Application Name	Name of the application.
Active Processes	The number of processes that are available.
Alive Processes	The number of live processes that are in use.
% Memory Utilization	The percentage of memory utilization.
%CPU Utilization	The percentage of CPU utilization.
I/O Rate (KB per Second)	The number of characters (in KB) per second transferred for processes in this group to all devices during the interval. This includes IO to disk, terminal, tape and printers

System Events- Shows a list of all the events generated on the system with their severity.

Parameters of the System Events

Name	Description
Event ID	Unique identifier for the event.
Severity	The list for UNIX is: Alert, Emergency, Error, Warning, Information, or Verbose
Time Stamp	The time when the event was generated
Description	Details of the event

The following options are available in the tabs:

Sorting

Shows the values in decreasing or increasing sequence. You can view the results as per the requirement. You can find out the process where the CPU utilization is highest. This option is available in all the tabs and in each column. Click (Clear Sorting) to remove the sort.

Search

The page contains a search box to filter data based on your requirement.

Options

The **Options** icon provides you the way to enable settings in the page. You can change the settings available on the page.

Field	Description
Real Time Settings	Contains the option to control the rate at which the graphs will refresh on the page. You can select the rate from the interval 1-10 seconds.
Refresh	Click to restart the graphs. This clears the previous results and you can view the current graphs.
Play	Click to start the real time graphs. By default, the real time graphs are available.
Pause	Click to freeze the page for sometime. You must click play to restart the real time graphs.

Field	Description
Stop	Click to halt the refresh of the real time graphs. You can analyze the reason of the utilization by using this icon as the graphs will not get refreshed.

You can drill down to check the details about the processes that are running, CPU usage, free disk space, swap memory, filesystem, Network usage, and Application running on VM. You can sort and view the results by each column in the available tabs. Additionally, you can search for the specific process ID or process name.

Use Case: Using Guest OS Drilldown to Troubleshoot the Resource Utilization in your Environment

The following section shows you a scenario where you can troubleshoot and find out the reason for the unexpected performance of the selected VM.

Scenario

Linda is a VMware Administrator and is monitoring the performance of all the Virtual Machines from the Treemap view. She observes that one of the VMs is red. To troubleshoot the reason of VM becoming red, she uses HPE Cloud Optimizer Guest OS drill down tool to troubleshoot the reason for the same.

The steps include the following:

1. Log on to HPE Cloud Optimizer user interface.
2. Go to the **Treemap view > VM Compute** to view the performance of all the Virtual Machines from the Treemap view.
3. Select the VM that is showing red or yellow color. Right-click on the selected VM and select Guest OS Drilldown option. The Guest OS page opens. You can view the resource utilization graphs. From the graphs, you can derive if the system is significantly using the memory or CPU.

You can additionally troubleshoot the cause of the performance by checking the following:

- View the **Process Cloud** tab to check which processes are using high CPU or memory.
- Go to the **Processes** tab to derive the resource utilization by the processes that are significantly highlighted in the **Process Cloud** tab.

- From the **CPU** tab, you can view the resource utilization for a particular CPU.
- Use the Disk, File System, Network, and Application tabs to view the utilization.

Chapter 11: HPE Cloud Optimizer API Reference

HPE Cloud Optimizer Representational State Transfer (REST) Application Programming Interfaces (APIs) provide a REST based interface to get Virtual Machine (VM) placement suggestions, optimization recommendations for your virtualized environment, view powered off and idle VMs, view resource allocation for VMs, and view snapshots.

The VM placement suggestions provide the most suitable Clusters or Hosts or Datastores in which the VM can be placed based on the Days To Capacity¹ for all the container entities. The optimization recommendation APIs provide information on VM sizing and datastore sizing. These recommendations can be used to automate reclamation of resources not being utilized, for example datastore space reclamation and for right sizing of VMs.

This document provides details about the REST APIs resource model schema and JavaScript Object Notation (JSON) examples.

HPE Cloud Optimizer provides the following set of REST APIs:

- [REST APIs for Placement](#)
- [REST APIs for Optimization Recommendations](#)
- [REST APIs to view Powered Off VMs](#)
- [REST APIs to view Idle VMs](#)
- [Using REST APIs to view CPU, Memory, and Storage allocation of VMs](#)
- [REST APIs to view Snapshots](#)

Using REST API for Placement

The HPE Cloud Optimizer REST API for Placement can be used by applications to get placement suggestions from HPE Cloud Optimizer whenever a new VM is created. The REST API for Placement provides placement suggestions for VMs in VMware, Hyper-V, and KVM domains only.

¹the number of days in which the entity reaches the maximum capacity

The REST API takes VM's specification, container's specification, and storage specification (optional) as input and provides suggestions on placing a VM in the most suitable cluster or host and datastore as an output. The REST API provides placement suggestions only if all the requested VMs can be placed in the container entities.

Resource Information

Response formats	JSON
Request Header requirements	Content-Type: application/json Authorization: Basic <Base 64 encoded username:password> (optional, required only when LDAP is enabled)
Supported request methods	POST

Resource URL

`https://<IP address or host name>:<port number>/PV/api/v1/placement`

In this instance,

- `<IP address or host name>` is the IP address or host name of the HPE Cloud Optimizer server.
- `<port number>` is the port number of the HPE Cloud Optimizer server. The default port number of HPE Cloud Optimizer server is 8444.

POST Parameters

Following is a representation of REST API request body:

Request Body [2]

outputFormat : flatlist

Optional input parameter for detailed list output.

vmSpecifications[n] *

0 {8}

numVM : 5

numCPU : 2

memSizeGB : 2

diskSizeGB : 10

expectedCPUUtil : 80

expectedMemUtil : 80

startDayFromToday : 1

endDayFromToday : 90

1 {8}

.

.

n {8}

containerSpecifications [m] *

0 {3}

virtType : VMWARE

type : VCENTER

vcenterName : 16.184.46.174

1 {3}

.

.

m {3}

storageSpecifications [x] *

0 {4}

vcenterName : examplevcenter.com

datacenterName : ispi_qa_dc_1029

clusterName : ispi_qa_cluster_1029

datastoreName : ovpesx9: storage1

* n is the number of VM specifications.

* m is the number of container specifications.

* x is the number of storage specifications.

* storageSpecifications is optional.

HPE Cloud Optimizer (3.00)

Page 250 of 337

The request body consists of the following objects:

- **outputFormat : flatlist:** This is an optional input parameter. It is required when you need the placement output in the form of a detailed list. If you do not specify any parameter, the output appears in form of a three-dimensional matrix.
- **vmSpecifications:** vmSpecifications is an array of JSON objects. The objects are used to specify the configurations of VMs to be placed. It has the following parameters:

Parameter	Description
numVM (<i>optional</i>)	The number of VMs of a specific configuration that are to be added. The default value is 1.
numCPU (<i>optional</i>)	The number of CPUs in each VM. The default value is 1.
memSizeGB	The memory size of each VM (in GigaBytes)
diskSizeGB	The disk size of each VM (in GigaBytes)
expectedCPUUtil (<i>optional</i>)	The expected CPU usage for a VM (in percentage). The default value is 100.
expectedMemUtil (<i>optional</i>)	The expected Memory usage for a VM (in percentage). The default value is 100.
startDayFromToday (<i>optional</i>)	The number of days from the current date on which you would like to create the VM. The default value is 1. For example, if the value is 5, the placement suggestions would be given considering that the VM has to be placed five days after the current date.
endDayFromToday (<i>optional</i>)	The number of days from the current date till which the VM would be active. The default value is 90. For example, if the value is 20, the placement suggestions would be given considering that the VM would remain active for 20 days from the current date.

- **containerSpecifications:** containerSpecifications is an array of JSON objects. The objects are used to specify the container configurations in which the VMs can be placed. It has the following parameters:

Parameter	Description
virtType	The type of the virtualized environment. Currently, placement suggestions are supported only on VMware,

Parameter	Description
	Hyper-V, and KVM virtualized environment. Hence, the value can be VMWARE, Hyper-V, or KVM.
type	<p>The type of container. You can specify type as VCENTER, DATACENTER, CLUSTER, HOST, or GROUP. This is a mandatory field.</p> <ul style="list-style-type: none"> • If the type is specified as VCENTER, the parameter vcenterName is mandatory. • If the type is specified as DATACENTER, the parameters vcenterName and datacenterName are mandatory. • If the type is specified as CLUSTER, the parameters vcenterName, datacenterName, and clusterName are mandatory. • If the type is specified as HOST, the hostName is mandatory. • If the type is specified as GROUP, the parameters groupType and groupName are mandatory. virtType is not required for a GROUP.
vcenterName	The name of the vCenter.
datacenterName	The name of the datacenter.
clusterName	The name of the cluster.
hostName	The name of the host.
groupName	The name of the group.
groupType	The type of the group.

Note: If the container type is GROUP, HPE Cloud Optimizer recommends that you specify a Business Group of standalone hosts.

- **storageSpecifications**(*optional*): storageSpecifications is an array of JSON objects. The objects are used to specify the datastores in which the VMs can be placed. If you do not specify this parameter, the placement suggestions are given considering all the datastores in the vCenter.

For individual datastore, the storageType is STORAGE and for a group of datastores, the storageType is STORAGE_GROUP. The storageType is a mandatory field.

When the storageType is STORAGE, the parameters required are as follows:

Parameter	Description
vcenterName	The name of the vCenter. It should be one of the vCenter names specified in the containerSpecifications.
datacenterName	The name of the datacenter.
clusterName	The name of the cluster.
datastoreName	The name of the datastore.

When the storageType is STORAGE_GROUP, the parameters required are as follows:

Parameter	Description
groupName	The name of the group.
groupType	The type of the group.

Custom Error Codes

If the REST API for Placement is unable to provide placement suggestions, it returns the following custom error codes:

Error Code	Error Description
512	Unable to provide placement suggestions as none of the available clusters and storage meet the input specifications.
513	Unable to provide placement suggestions due to lack of forecast information.
514	Unable to provide placement suggestions as there are not enough resources available to place the requested VMs.

Sample API

URL

```
https://<IP address or host name>:<port number>/PV/api/v1/placement
```

The request body is as follows:

```
{  
  "vmSpecifications": [  
    {
```

```
        "numVM":5,
        "numCPU":2,
        "memSizeGB":2,
        "diskSizeGB":10,
        "expectedCPUUtil":80,
        "expectedMemUtil":80,
        "startDayFromToday":1,
        "endDayFromToday":90
    },
    {
        "numVM":4,
        "numCPU":4,
        "memSizeGB":8,
        "diskSizeGB":5,
        "expectedCPUUtil":20,
        "expectedMemUtil":40,
        "startDayFromToday":1,
        "endDayFromToday":90
    }
],
"containerSpecifications":[
{
    "virtType":"VMWARE",
    "type":"VCENTER",
    "vcenterName":"examplevcenter.com",
    "datacenterName" : "ispi_qa_DC_1029",
    "clusterName" : "ispi_qa_CLUSTER_1029"
},
```

```
{
    "virtType":"VMWARE",
    "type":"DATACENTER",
    "vcenterName":"16.184.45.94",
    "datacenterName" : "DataCenter1",
    "clusterName" : "PMi_Cluster"
},
{
    "virtType":"VMWARE",
    "type":"CLUSTER",
    "vcenterName":"examplevcenter.com",
    "datacenterName" : "ispi_qa_DC_1030",
    "clusterName" : "ispi_qa_CLUSTER_1030"
},
```

```
{
    "virtType" : "VMWARE",
    "type" : "HOST",
    "hostName" : "16.184.42.73"
},
{
    "type" : "GROUP",
    "groupType" : "Departments",
    "groupName" : "Finance",
}
],
```

```
"storageSpecifications":[
{
```

```
    "storageType" : "STORAGE"
    "vcenterName" : "examplevcenter.com",
    "datacenterName" : "ispi_qa_DC_1029",
    "clusterName" : "ispi_qa_CLUSTER_1029",
    "datastoreName" : "ovpesx9:storage1"
  },
  {
    "storageType" : "STORAGE"
    "vcenterName" : "16.184.45.94",
    "datacenterName" : "DataCenter1",
    "clusterName" : "PMi_Cluster",
    "datastoreName" : "datastore1_1"
  },
  {
    "storageType" : "STORAGE"
    "vcenterName" : "examplevcenter.com",
    "datacenterName" : " ispi_qa_DC_1030",
    "clusterName" : " ispi_qa_DC_1030",
    "datastoreName" : "storage3"
  },
  {
    "storageType" : "STORAGE"
    "vcenterName" : "examplevcenter.com",
    "datacenterName" : " ispi_qa_DC_1030",
    "clusterName" : " ispi_qa_DC_1030",
    "datastoreName" : "storage4"
  },
}
```



```
{
    "storageType" : "STORAGE_GROUP"
    "groupType" : "DepartmentsStorage",
    "groupName" : "FinanceStorage"
}
```

JSON Response

The JSON Response by the REST API can be of two different types: Detailed List Output and Three-Dimensional List Output.

- Detailed List Output: To get the JSON response in the form of a Detailed List Output, specify **"outputFormat" : "flatlist"** in the request body.
- Three-Dimensional List Output: To get the JSON response in the form of a Three-Dimensional List Output, do not specify **"outputFormat" : "flatlist"** in the request body.

Following is a brief explanation of the two different types of JSON response given by the REST API:

Detailed List Output

The JSON response contains an array of JSON objects. Each JSON object corresponds to a VM to be placed. Each JSON object contains the VM configuration details, container and the datastore details where the VM can be placed.

The JSON output can be visually represented in the following way:



```

Result [9]
  0 {3}
  1 {3}
    vmDetails {7}
      numCPU : 2
      memSizeGB : 2
      diskSizeGB : 10
      expectedCPUUtil : 80
      expectedMemUtil : 80
      startDayFromToday : 1
      endDayFromToday : 90
    containerDetails {5}
      virtType : VMWARE
      type : CLUSTER
      vcenterName : examplevcenter.com
      datacenterName : ispi_qa_DC_1030
      clusterName : ispi_qa_CLUSTER_1030
    storageDetails {4}
      vcenterName : examplevcenter.com
      datacenterName : ispi_qa_DC_1030
      clusterName : ispi_qa_DC_1030
      datastoreName : storage3
  2 {3}
  3 {3}
  4 {3}
  5 {3}
  6 {3}
  7 {3}
  8 {3}
  
```

Configuration details, container details, and storage details of the VM to be placed.

Three-Dimensional Matrix Output

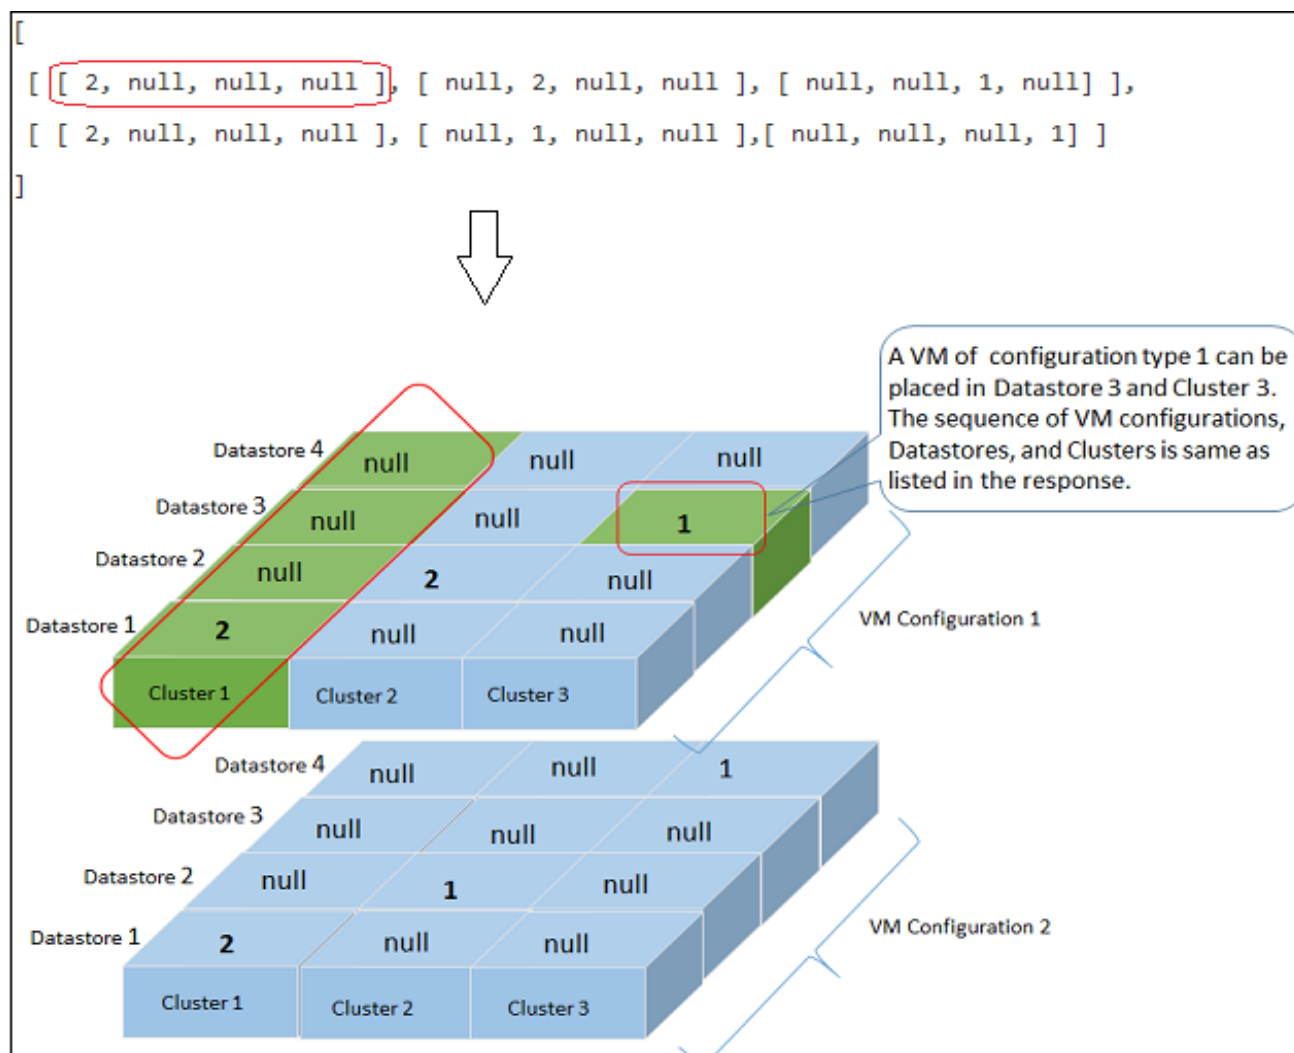
The API provides placement suggestions in the form of matrix. The matrix can be interpreted after understanding other output parameters. Following is a brief outline of the output parameters:

Parameter	Description
numVMConfigs	The total number of VM configurations.
numContainers	The total number of distinct clusters in which the VMs can be placed.

Parameter	Description
numDatastores	The total number of distinct datastores in which the VMs can be placed.
vmConfigDetailsList	This object contains the list of VM configurations. Note the sequence of VM configurations in the list as it would be required to interpret the result matrix.
containerDetailsList	This object contains the list of containers in which the VMs can be placed. Note the sequence of containers in the list as it would be required to interpret the result matrix.
storageDetailsList	This object contains the list of datastores in which the VMs can be placed. Note the sequence of datastores in the list as it would be required to interpret the result matrix.

resultMatrix: The placement suggestions are given in the form of a three-dimensional matrix. The result matrix in the sample output is:

```
[  
  [ [ 2, null, null, null ], [ null, 2, null, null ], [ null, null, 1, null]  
],  
  [ [ 2, null, null, null ], [ null, 1, null, null ], [ null, null, null, 1 ] ]  
]
```



The following table gives the container details and corresponding output specification level:

Container Details	Output Specification Level
VCENTER, CLUSTER, DATACENTER	CLUSTER
HOST, GROUP	HOST
Mix of (VCENTER, CLUSTER, DATACENTER) and (HOST, GROUP)	HOST

Sample Output Details

Example 1 : When container includes VCENTER, CLUSTER, or DATACENTER.

Detailed List Output

The following JSON response appears if you specify **"outputFormat" : "flatlist"** in the request body:

```
"result" : [  
{  
  "vmDetails" : {  
    "numCPU":2,  
    "memSizeGB":2,  
    "diskSizeGB":10,  
    "expectedCPUUtil":80,  
    "expectedMemUtil":80,  
    "startDayFromToday":1,  
    "endDayFromToday":90  
  },  
  "containerDetails" : {  
    "virtType":"VMWARE",  
    "type":"CLUSTER",  
    "vcenterName":"examplevcenter.com",  
    "datacenterName" : "ispi_qa_DC_1030",  
    "clusterName" : "ispi_qa_CLUSTER_1030"  
  },  
  "storageDetails" : {  
    "vcenterName" : "examplevcenter.com",  
    "datacenterName" : " ispi_qa_DC_1030",  
    "clusterName" : " ispi_qa_DC_1030",  
    "datastoreName" : "storage3"  
  }  
},  
  "vmDetails" : {  
    "numCPU":2,
```

```
        "memSizeGB":2,
        "diskSizeGB":10,
        "expectedCPUUtil":80,
        "expectedMemUtil":80,
        "startDayFromToday":1,
        "endDayFromToday":90
    },
    "containerDetails" : {
        "virtType":"VMWARE",
        "type":"VCENTER",
        "vcenterName":"examplevcenter.com",
        "datacenterName" : "ispi_qa_DC_1029",
        "clusterName" : "ispi_qa_CLUSTER_1029"
    },
    "storageDetails" : {
        "vcenterName" : "examplevcenter.com",
        "datacenterName" : "ispi_qa_DC_1029",
        "clusterName" : "ispi_qa_CLUSTER_1029",
        "datastoreName" : "ovpesx9:storage1"
    }
},
.....
}
],
"daysToCapacity" : 30,
"errorCode" : 0
}
```

Three-Dimensional Matrix Output

The following JSON response appears if you do not specify **"outputFormat" : "flatlist"** in the request body.

```
{
  "numVMConfigDetails" : 2,
  "numContainerDetails" : 3,
  "numStorageDetails" : 4,
  "vmConfigDetailsList" : [
    {
      "numVM":5,
      "numCPU":2,
      "memSizeGB":2,
      "diskSizeGB":10,
      "expectedCPUUtil":80,
      "expectedMemUtil":80,
      "startDayFromToday":1,
      "endDayFromToday":90
    },
    {
      "numVM":4,
      "numCPU":4,
      "memSizeGB":8,
      "diskSizeGB":5,
      "expectedCPUUtil":20,
      "expectedMemUtil":40,
      "startDayFromToday":1,
      "endDayFromToday":90
    } ],
  "containerDetailsList" : [ {
    "virtType" : "VMWARE",
```

```
        "type" : "CLUSTER",
        "vcenterName" : "examplevcenter.com",
        "datacenterName" : "ispi_qa_DC_1029",
        "clusterName" : "ispi_qa_CLUSTER_1029",
        "daysToCapacity" : 30
    },
    {
        "virtType" : "VMWARE",
        "type" : "CLUSTER",
        "vcenterName" : "16.184.45.94",
        "datacenterName" : "DataCenter1",
        "clusterName" : "PMi_Cluster",
        "daysToCapacity" : 70
    },
    {
        "virtType" : "VMWARE",
        "type" : "CLUSTER",
        "vcenterName" : "examplevcenter.com",
        "datacenterName" : " ispi_qa_DC_1030",
        "clusterName" : " ispi_qa_DC_1030",
        "daysToCapacity" : 50}
],
"storageDetailsList" : [
{
    "vcenterName" : "examplevcenter.com",
    "datacenterName" : "ispi_qa_DC_1029",
    "clusterName" : "ispi_qa_CLUSTER_1029",
    "datastoreName" : "ovpesx9:storage1",
```



```
        "daysToCapacity" : 90
    },
    {
        "vcenterName" : "16.184.45.94",
        "datacenterName" : "DataCenter1",
        "clusterName" : "PMi_Cluster",
        "datastoreName" : "datastore1_1",
        "daysToCapacity" : 90
    },
    {
        "vcenterName" : "examplevcenter.com",
        "datacenterName" : " ispi_qa_DC_1030",
        "clusterName" : " ispi_qa_DC_1030",
        "datastoreName" : "storage3",
        "daysToCapacity" : 40
    },
    {
        "vcenterName" : "examplevcenter.com",
        "datacenterName" : " ispi_qa_DC_1030",
        "clusterName" : " ispi_qa_DC_1030",
        "datastoreName" : "storage4",
        "daysToCapacity" : 60
    }
],
"resultMatrix" : [ [ [ 2, null, null, null ], [ null, 2, null, null], [ null,
null, 1, null] ], [ [ 2, null, null, null ], [ 1, null, null, null ],[ 1,
null, null, null ] ] ],
"daysToCapacity" : 30
```

Example 2 : When containers includes VCENTER, CLUSTER, DATACENTER, HOST, or GROUP.

Detailed List Output

```
{
  [ {
    "vmDetails" : {
      "numCPU" : 2,
      "memSizeGB" : 2.0,
      "diskSizeGB" : 10.0,
      "startDayFromToday" : 1,
      "endDayFromToday" : 90
    },
    "containerDetails" : {
      "type" : "HOST",
      "hostName" : "16.184.40.68",
      "daysToCapacity" : 90,
      "instanceID" : 43
    },
    "storageDetails" : {
      "vcenterName" : "16.184.45.37",
      "datacenterName" : "Gringotts",
      "clusterName" : "Personal",
      "datastoreName" : "datastore1 (10)",
      "daysToCapacity" : 90,
      "instanceID" : 16
    }
  }, {
    "vmDetails" : {
      "numCPU" : 4,
```

```
        "memSizeGB" : 8.0,
        "diskSizeGB" : 5.0,
        "startDayFromToday" : 1,
        "endDayFromToday" : 90
    },
    "containerDetails" : {
        "type" : "HOST",
        "hostName" : "16.184.40.118",
        "daysToCapacity" : 90,
        "instanceID" : 44
    },
    "storageDetails" : {
        "vcenterName" : "16.184.45.37",
        "datacenterName" : "Gringotts",
        "clusterName" : "Personal",
        "datastoreName" : "datastore1 (7)",
        "daysToCapacity" : 90,
        "instanceID" : 15
    }
},
...
...
...
{
} ],
        "daysToCapacity" : 90,
        "errorCode" : 0
    }
}
```

Three-Dimensional Matrix Output

The following JSON response appears if you do not specify **"outputFormat" : "flatlist"** in the request body.

```
{
  "numVMConfigDetails" : 2,
  "numContainerDetails" : 3,
  "numStorageDetails" : 3,
  "vmConfigDetailsList" : [ {
    "numVM" : 5,
    "numCPU" : 2,
    "memSizeGB" : 2.0,
    "diskSizeGB" : 10.0,
    "expectedCPUUtil" : 80.0,
    "expectedMemUtil" : 80.0,
    "startDayFromToday" : 1,
    "endDayFromToday" : 90
  },
  {
    "numVM" : 4,
    "numCPU" : 4,
    "memSizeGB" : 8.0,
    "diskSizeGB" : 5.0,
    "expectedCPUUtil" : 20.0,
    "expectedMemUtil" : 40.0,
    "startDayFromToday" : 1,
    "endDayFromToday" : 90
  }
],
  "containerDetailsList" : [ {
```

```
        "type" : "HOST",
        "hostName" : "16.184.40.68",
        "daysToCapacity" : 90,
        "instanceID" : 43
    },
    {
        "type" : "HOST",
        "hostName" : "16.184.40.118",
        "daysToCapacity" : 90,
        "instanceID" : 44
    },
    {
        "type" : "HOST",
        "hostName" : "16.184.40.123",
        "daysToCapacity" : 90,
        "instanceID" : 45
    }
],
    "storageDetailsList" : [ {
        "vcenterName" : "16.184.45.37",
        "datacenterName" : "Gringotts",
        "clusterName" : "Business",
        "datastoreName" : "datastore1 (8)",
        "daysToCapacity" : 90,
        "instanceID" : 10
    },
    {
        "vcenterName" : "16.184.45.37",
```

```
        "datacenterName" : "Gringotts",
        "clusterName" : "Personal",
        "datastoreName" : "datastore1 (7)",
        "daysToCapacity" : 90,
        "instanceID" : 15
    },
    {
        "vcenterName" : "16.184.45.37",
        "datacenterName" : "Gringotts",
        "clusterName" : "Personal",
        "datastoreName" : "datastore1 (10)",
        "daysToCapacity" : 90,
        "instanceID" : 16
    }
],
"resultMatrix" : [ [ [ null, null, 5 ], [ null, null, null ], [ null, null, null ] ], [ [ null, null, null ], [ null, 3, null ], [ 1, null, null ] ] ],
"daysToCapacity" : 90,
"errorCode" : 0
}
```

Using REST APIs for Optimization Recommendations

The HPE Cloud Optimizer REST APIs for Optimization Recommendations provides sizing information for VMs and Datastores in Datacenters, Clusters, and Hosts in your virtualized infrastructure.

Resource Model

REST APIs for optimization recommendations can be invoked at two levels:

- Collection of Entities

At this level, you can view information for collection of Datacenters, Clusters, Hosts, and Datastores in your virtualized infrastructure.

- Specific Entities

At this level, you can view information for specific instances of Datacenters, Clusters, Hosts, Datastores, and VMs. You can view VM sizing and datastore sizing information for specific Datacenter, Cluster, and Host instances only.

Note: You may observe a difference in the count of VMs in the Treemap and the count of VMs returned by the REST API. This is because the REST API considers only the VMs in the virtualized infrastructure whereas the Treemap includes the VM Templates as well.

Resource Information

Response formats	JSON
Request Header requirements	Content-Type: application/json Authorization: Basic <Base 64 encoded username:password> (optional, required only when LDAP is enabled)
Supported request methods	GET

Resource URL

The Base URL to get information about the entities is:

`https://<ip_address or host_name>:<port_number>/PV/api/v1`

In this instance,

- `ip_address` or `host_name` is the IP address or host name of the HPE Cloud Optimizer server.
- `port_number` is the port number of the HPE Cloud Optimizer server. The default port number is 8444.

The Base URL must be appended with appropriate URIs to get information about various entities monitored by HPE Cloud Optimizer server. The following table summarizes REST API URL details for various entities:

Click on the URL to view specific examples.

Entity	API	URL
VM	To get	<a href="https://<ip_address>:<port_">https://<ip_address>:<port_

Entity	API	URL
	information of all the VMs along with sizing details.	<code>number>/PV/api/v1/vm</code>
	To get information of a specific VM.	<code>https://<ip_address>:<port_number>/PV/api/v1/vm/<instance_id></code> In this instance, instance_id is the resource ID of the VM.
Datastore	To get datastore information along with sizing details.	<code>https://<ip_address>:<port_number>/PV/api/v1/datastore</code>
	To get information of a specific datastore.	<code>https://<ip_address>:<port_number>/PV/api/v1/datastore/<instance_id></code> In this instance, instance_id is the resource ID of the datastore.
	To get information about all the VMs on a specific datastore.	<code>https://<ip_address>:<port_number>/PV/api/v1/datastore/<instance_id>/vm</code> In this instance, instance_id is the resource ID of the datastore.
	To get information of a specific VM in a datastore that is filtered on the basis of Status.	<code>https://<ip_address>:<port_number>/PV/api/v1/datastore/<instance_id>/vm?status=<vm_status></code> In this instance, instance_id is the resource ID of the datastore. status is the current status of the VM. The parameter status can have the values as idle, mostlypoweredoff, deleted, or active.

Entity	API	URL
Cluster	To get information of all clusters.	<a href="https://<ip_address>:<port_number>/PV/api/v1/cluster">https://<ip_address>:<port_number>/PV/api/v1/cluster
	To get information of a specific cluster.	<a href="https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>">https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id> In this instance, instance_id is the resource ID of the cluster.
	To get sizing information of all the VMs in a specific cluster.	<a href="https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>/vmsizing">https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>/vmsizing In this instance, instance_id is the resource ID of the cluster.
	To get sizing information of a specific VM in a cluster that is filtered based on name and sizing parameters.	<a href="https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>/vmsizing?name=<system_name>&sizing=<sizing>">https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>/vmsizing?name=<system_name>&sizing=<sizing> In this instance, instance_id is the resource ID of the Cluster. system_name is the VM display name. You can specify a sub-string of the VM display name. sizing is the CPU or Memory size of the specific VM. The parameter sizing can have the values as oversized, undersized, or rightsized.
	To get sizing information of all the datastores that are part of a specific cluster.	<a href="https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>/datastoresizing">https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>/datastoresizing In this instance, instance_id is the resource ID of the cluster.
Datcenter	To get	<a href="https://<ip_address>:<port_number>">https://<ip_address>:<port_number>

Entity	API	URL
	information of all the Datacenters.	number>/PV/api/v1/datacenter
	To get information of a specific Datacenter	https://<ip_address>:<port_number>/PV/api/v1/datacenter/<instance_id> In this instance, instance_id is the resource ID of the Datacenter.
	To get sizing information of all the VMs that are part of a specific Datacenter.	https://<ip_address>:<port_number>/PV/api/v1/datacenter/<instance_id>/vmsizing In this instance, instance_id is the resource ID of the Datacenter.
	To get sizing information of a specific VM in a Datacenter that is filtered based on name and sizing parameters.	https://<ip_address>:<port_number>/PV/api/v1/datacenter/<instance_id>/vmsizing?name=<system_name>&sizing=<sizing> In this instance, instance_id is the resource ID of the Datacenter. system_name is the VM display name. You can specify a sub-string of the VM display name. sizing is the CPU or Memory size of the specific VM. The parameter sizing can have the values as oversized, undersized, or rightsized.
	To get sizing information of all the datastores that are part of a specific Datacenter	https://<ip_address>:<port_number>/PV/api/v1/datacenter/<instance_id>/datastoresizing In this instance, instance_id is the resource ID of the Datacenter.
Host	To get	https://<ip_address>:<port_

Entity	API	URL
	information of all the Hosts.	<code>number>/PV/api/v1/host</code>
	To get information of a specific Host.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/host/<instance_id></code></p> <p>In this instance, instance_id is the resource ID of the Host.</p>
	To get sizing information of all the VMs in a specific Host.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/host/<instance_id>/vmsizing</code></p> <p>In this instance, instance_id is the resource ID of the Host.</p>
	To get sizing information of a specific VM in a Host that is filtered based on name and sizing parameters.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/host/<instance_id>/vmsizing?name=<system_name>&sizing=<sizing></code></p> <p>In this instance, instance_id is the resource ID of the Host. system_name is the VM display name. You can specify a sub-string of the VM display name. sizing is the CPU or Memory size of the specific VM. The parameter sizing can have the values as oversized, undersized, or rightsized.</p>
	To get datastore information and sizing details of all the datastores in a specific Host.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/host/<instance_id>/datastoresizing</code></p> <p>In this instance, instance_id is the resource ID of the Host.</p>

Datacenter

The following section provides sample JSON response generated by the REST API for optimization recommendations for Datacenter.

To get information of all the Datacenters

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datacenter
```

The following JSON is returned in the response:

```
{
  "instancecount" : 3,
  "instances" : [
    {
      "properties" :
      {
        "ParentUUID" : "16.184.45.37",
        "SystemVirtType" : "VMWARE",
        "SystemID" : "16.184.45.37/SMV-BLADE",
        "ParentType" : "VCENTER",
        "SystemRole" : "DATACENTER",
        "LSName" : "16.184.45.37/SMV-BLADE",
        "SystemName" : "SMV-BLADE"
      },
      "resourcepath" : "/datacenter/1",
      "resourceid" : 1
    },
    {
      "properties" : {
        "ParentUUID" : "16.184.45.37",
```

```
        "SystemVirtType" : "VMWARE",
        "SystemID" : "16.184.45.37/Google.Cloud",
        "ParentType" : "VCENTER",
        "SystemRole" : "DATACENTER",
        "LSName" : "16.184.45.37/Google.Cloud",
        "SystemName" : "Google.Cloud"
    },
    "resourcepath" : "/datacenter/2",
    "resourceid" : 2
}....
```

To get information of a specific Datacenter

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datacenter/<instance_
id>
```

The following JSON is returned in the response:

```
{
    "instance" : [
        {
            "properties" :
            {
                "ParentUUID" : "16.184.45.37",
                "SystemVirtType" : "VMWARE",
                "SystemID" : "16.184.45.37/Google.Cloud",
                "ParentType" : "VCENTER",
                "SystemRole" : "DATACENTER",
                "LSName" : "16.184.45.37/Google.Cloud",
                "SystemName" : "Google.Cloud"
            }
        }
    ]
}
```

```
    },  
    "resourcepath" : "/datacenter/2",  
    "resourceid" : 2  
  } ]  
}
```

To get sizing information of all the VMs that are part of a specific Datacenter

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datacenter/<instance_  
id>/vmsizing
```

The following JSON is returned in the response:

```
{  
  "instancecount" : 24,  
  "instances" : [  
    {  
      "properties" :  
        {  
          "SystemOSName" : "Red Hat Enterprise Linux 6 (64-bit)",  
          "StatTime" : "",  
          "SystemID" : "421ad522-d7f5-0025-1b0a-409df1a2898f",  
          "SystemOSRelease" : "",  
          "ParentType" : "HOST",  
          "SystemPath" : "[BLR-LUN-0112]  
martellvm40158.hpswlab.hp.com/martellvm40158.hpswlab.hp.com.vmx",  
          "SystemRole" : "GUEST",  
          "Type" : "",  
          "UseHardwareAssistedVirtualization" : "",  
          "MacAddress" : "00:50:56:9a:79:aa",
```

```
"SystemHostHostName" : "agentesx.ind.hp.com",  
"SystemVirtType" : "VMWARE",  
"ClusterName" : "BLR-APPLE-CLUSTER",  
"SystemOSVersion" : "",  
"ManagementIP" : "",  
"FlavorId" : "",  
"VCIPAddress" : "",  
"SystemVirtPlatform" : "",  
"InstanceName" : "",  
"CPUArch" : "",  
"ProcessorVendor" : "",  
"VirtVersion" : "",  
"SystemName" : "martellvm40158.hpswlab.hp.com",  
"MarkedForDeletion" : "",  
"VMOwnerId" : "",  
"HPCSPort" : "",  
"LSID" : "79",  
"HAEnabled" : "",  
"ParentUUID" : "38393636-3430-4753-4832-333442364c53",  
"DynamicMemoryEnabled" : "",  
"LSMode" : "UnCapped",  
"CPUVendor" : "",  
"HasSnapshot" : "",  
"BelongsToDatacenter" : "16.184.45.37/BLR-APPLE-DC",  
"VMHost_UUID" : "",  
"LSName" : "421ad522-d7f5-0025-1b0a-409df1a2898f",  
"VirtVersionState" : "",  
"ManagementUUID" : "",
```

```
        "LSShared" : "Shared",
        "HPCSProtocol" : "",
        "SystemHostName" : "ci-0050569A79AA"
    },
    "resourcepath" : "/vm/62",
    "resourceid" : 62,
    "sizing" :
    {
        "PO_MemCapacity" : 10.0,
        "PO_CPUCycleEntlMin" : 0.0,
        "PO_RecommendedNumCPU" : 1.0,
        "PO_RecommendedCPUReservation" : 0.0,
        "PO_RecommendedMem" : 5.0,
        "PO_MemEntlMin" : 0.0,
        "PO_RecommendedMemReservation" : 0.0,
        "PO_NumCpu" : 2
    },
    "MemSize" : "oversized",
    "CPUSize" : "oversized"
}, .....
```

To get sizing information of a specific VM in a Datacenter filtered based on name and sizing parameters

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datacenter/<instance_id>/vmsizing?name=<system_name>&sizing=<sizing>
```

```
{
    "instancecount" : 1,
```



```
"instances" : [ {  
  "properties" : {  
    "SystemOSName" : "CentOS 4/5/6 (64-bit)",  
    "StatTime" : "",  
    "SystemID" : "420c9ec7-a54c-7173-ae2d-a6daa3621419/ispiva1.ind.hp.com_  
vPV1.2_MR",  
    "SystemPath" : "[Storage2]ispiva1.ind.hp.com_vPV1.2_MR/ispiva1.ind.hp.com_  
vPV1.2_MR.vmx",  
    "SystemOSRelease" : "",  
    "ParentType" : "HOST",  
    "Type" : "",  
    "SystemRole" : "GUEST",  
    "UseHardwareAssistedVirtualization" : "",  
    "SystemHostHostName" : "ovpesx9.ind.hp.com",  
    "MacAddress" : "00:50:56:8c:11:f6,52:54:00:24:e2:1e",  
    "SystemVirtType" : "VMWARE",  
    "ClusterName" : "ispi_qa_CLUSTER_1029",  
    "SystemOSVersion" : "",  
    "ManagementIP" : "",  
    "FlavorId" : "",  
    "VCIPAddress" : "",  
    "SystemVirtPlatform" : "",  
    "CPUArch" : "",  
    "InstanceName" : "",  
    "ProcessorVendor" : "",  
    "VirtVersion" : "",  
    "SystemName" : "ispiva1.ind.hp.com_vPV1.2_MR",  
    "MarkedForDeletion" : "",  
    "VMOwnerId" : "",
```

```
    "HPCSPort" : "",
    "LSID" : "1437",
    "HAEnabled" : "",
    "ParentUUID" : "34313734-3537-5347-4837-323034535242",
    "DynamicMemoryEnabled" : "",
    "LSMode" : "UnCapped",
    "CPUVendor" : "",
    "HasSnapshot" : "",
    "BelongsToDatacenter" : "iwfvm01029.hpswlab.s.adapps.hp.com/ispiva1.qa_DC_1029",
    "VMHost_UUID" : "",
    "LSName" : "420c9ec7-a54c-7173-ae2d-a6daa3621419/ispiva1.ind.hp.com_vPV1.2_
MR",
    "VirtVersionState" : "",
    "LSShared" : "Shared",
    "ManagementUUID" : "",
    "SystemHostName" : "ispiva1.ind.hp.com",
    "HPCSProtocol" : ""
  },
  "resourcepath" : "/vm/4",
  "resourceid" : 4,
  "sizing" : {
    "MemEntlMin" : 0.0,
    "RecommendedNumCPU" : 1.0,
    "RecommendedCPUReservation" : 0.0,
    "NumCpu" : 2,
    "CPUCycleEntlMin" : 0.0,
    "RecommendedMemReservation" : 0.0,
    "RecommendedMem" : 2.0,
```

```
"MemCapacity" : 4.0
},
"MemSize" : "oversized",
"CPUSize" : "oversized"
} ]
}
```

To get sizing information of all the datastores that are part of a specific Datacenter

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datacenter/<instance_
id>/datastoresizing
```

The following JSON is returned in the response:

```
{
  "instance" : [
    {
      "properties" :
      {
        "ParentUUID" : "16.184.45.37/BLR-APPLE-DC/BLR-APPLE-CLUSTER",
        "SystemVirtType" : "VMWARE",
        "SystemID" : "4fb1a0b7-9de543d2-b203-ac162d75e5fc",
        "ParentType" : "CLUSTER",
        "ClusterName" : "BLR-APPLE-CLUSTER",
        "BelongsToDatacenter" : "16.184.45.37/BLR-APPLE-DC",
        "Type" : "VMFS",
        "SystemRole" : "DATASTORE",
        "LSName" : "4fb1a0b7-9de543d2-b203-ac162d75e5fc",
        "SystemName" : "BLR-LUN-0112"
      },

```

```
"resourcepath" : "/datastore/8",  
"resourceid" : 8,  
"sizing" :  
{  
    "PO_StorageReclaimable" : 138.72,  
    "PO_StorageAllocationPossibleAfterReclaim" : 8973.99,  
    "PO_DiskUsage" : 1536.87  
}  
, .....
```

Cluster

The following section gives sample JSON responses generated by the REST API for optimization recommendations for Cluster.

To get information of all clusters

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/cluster
```

The following JSON is returned in the response:

```
{  
    "instancecount" : 4,  
    "instances" : [  
        {  
            "properties" :  
            {  
                "ParentUUID" : "",  
                "SystemVirtType" : "Hyper-V",  
                "IsFullyCached" : "",  
                "SystemID" : "hyperv-clus",  
                "BelongsToDatacenter" : "",
```

```
        "ParentType" : "",
        "Type" : "",
        "SystemRole" : "CLUSTER",
        "LSName" : "",
        "SystemName" : "hyperv-clus",
        "MarkedForDeletion" : "",
        "SystemVirtPlatform" : ""
    },
    "resourcepath" : "/cluster/1",
    "resourceid" : 1
}, ....
{
    "properties" :
    {
        "ParentUUID" : "16.184.45.37/SMV-BLADE",
        "SystemVirtType" : "VMWARE",
        "IsFullyCached" : "",
        "SystemID" : "16.184.45.37/SMV-BLADE/SMV-Cluster",
        "BelongsToDatacenter" : "16.184.45.37/SMV-BLADE",
        "ParentType" : "DATACENTER",
        "Type" : "VMWare DRS",
        "SystemRole" : "CLUSTER",
        "LSName" : "16.184.45.37/SMV-BLADE/SMV-Cluster",
        "SystemName" : "SMV-Cluster",
        "MarkedForDeletion" : "",
        "SystemVirtPlatform" : ""
    },
    "resourcepath" : "/cluster/2",
```

```
    "resourceid" : 2
}
```

To get information of a specific cluster

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>
```

The following JSON is returned in the response:

```
{
  "instance" : [
    {
      "properties" :
      {
        "ParentUUID" : "16.184.45.37/SMV-BLADE",
        "SystemVirtType" : "VMWARE",
        "IsFullyCached" : "",
        "SystemID" : "16.184.45.37/SMV-BLADE/SMV-Cluster",
        "BelongsToDatacenter" : "16.184.45.37/SMV-BLADE",
        "ParentType" : "DATACENTER",
        "Type" : "VMWare DRS",
        "SystemRole" : "CLUSTER",
        "LSName" : "16.184.45.37/SMV-BLADE/SMV-Cluster",
        "SystemName" : "SMV-Cluster",
        "MarkedForDeletion" : "",
        "SystemVirtPlatform" : ""
      },
      "resourcepath" : "/cluster/2",
      "resourceid" : 2
    }
  ]
}
```

```
]
}
```

To get sizing information of all the VMs in a specific cluster

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>/vmsizing
```

The following JSON is returned in the response:

```
{
  "instancecount" : 58,
  "instances" : [
    {
      "properties" :
      {
        "SystemOSName" : "",
        "StatTime" : "",
        "SystemID" : "1AC13360-8C92-47C7-BFD7-83CB4EA4A266",
        "SystemOSRelease" : "",
        "ParentType" : "HOST",
        "SystemPath" : "",
        "SystemRole" : "GUEST",
        "Type" : "",
        "UseHardwareAssistedVirtualization" : "",
        "MacAddress" : "00:15:5D:2F:78:49",
        "SystemHostHostName" : "GOLDRICH6.iwflabs.com",
        "SystemVirtType" : "Hyper-V",
        "ClusterName" : "hyperv-clus",
        "SystemOSVersion" : "",

```

```
    "ManagementIP" : "",
    "FlavorId" : "",
    "VCIPAddress" : "",
    "SystemVirtPlatform" : "",
    "InstanceName" : "",
    "CPUArch" : "",
    "ProcessorVendor" : "",
    "VirtVersion" : "",
    "SystemName" : "2008WINSOFT",
    "MarkedForDeletion" : "",
    "VMOwnerId" : "",
    "HPCSPort" : "",
    "LSID" : "",
    "HAEnabled" : "",
    "ParentUUID" : "35353636-3235-4753-4834-303241525843",
    "DynamicMemoryEnabled" : "False",
    "LSMode" : "",
    "CPUVendor" : "",
    "HasSnapshot" : "",
    "BelongsToDatacenter" : "",
    "VMHost_UUID" : "",
    "LSName" : "",
    "VirtVersionState" : "",
    "ManagementUUID" : "",
    "LSShared" : "",
    "HPCSProtocol" : "",
    "SystemHostName" : ""
  },
```



```
"resourcepath" : "/vm/3",  
"resourceid" : 3,  
"sizing" : { },  
"MemSize" : null,  
"CPUSize" : null  
}, ...
```

To get sizing information of a specific VM in a cluster filtered based on name and sizing parameters

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_id>/vmsizing?name=<system_name>&sizing=<sizing>
```

The following JSON is returned in the response:

```
{  
  "instancecount" : 1,  
  "instances" : [ {  
    "properties" : {  
      "SystemOSName" : "Microsoft Windows Server 2008 R2 (64-bit)",  
      "StatTime" : "",  
      "SystemID" : "421cc483-36a7-55a8-fc4f-fe3ee4715e9a/x86vm72 - OMW9 - abc",  
      "SystemPath" : "[IBTLUN21]x86vm72/x86vm72.vmx",  
      "SystemOSRelease" : "",  
      "ParentType" : "HOST",  
      "Type" : "",  
      "SystemRole" : "GUEST",  
      "UseHardwareAssistedVirtualization" : "",  
      "SystemHostHostName" : "ibtvz11.ind.hp.com",  
      "MacAddress" : "",
```

```
"SystemVirtType" : "VMWARE",
"ClusterName" : "Cluster01",
"SystemOSVersion" : "",
"ManagementIP" : "",
"FlavorId" : "",
"VCIPAddress" : "",
"SystemVirtPlatform" : "",
"CPUArch" : "",
"InstanceName" : "",
"ProcessorVendor" : "",
"VirtVersion" : "",
"SystemName" : "x86vm72 - OMW9 - Rajan",
"MarkedForDeletion" : "",
"VMOwnerId" : "",
"HPCSPort" : "",
"LSID" : "98",
"HAEnabled" : "",
"ParentUUID" : "30333735-3738-4753-4830-343758573459",
"DynamicMemoryEnabled" : "",
"LSMode" : "UnCapped",
"CPUVendor" : "",
"HasSnapshot" : "",
"BelongsToDatacenter" : "ibtv2.ind.hp.com/Datacenter01",
"VMHost_UUID" : "",
"LSName" : "421cc483-36a7-55a8-fc4f-fe3ee4715e9a/x86vm72 - OMW9 - abc",
"VirtVersionState" : "",
"LSShared" : "Shared",
"ManagementUUID" : "",
```

```
        "SystemHostName" : "x86vm72.ind.hp.com",
        "HPCSProtocol" : ""
    },
    "resourcepath" : "/vm/75",
    "resourceid" : 75,
    "sizing" : {
        "MemEntlMin" : 0.0,
        "RecommendedNumCPU" : 1.0,
        "RecommendedCPUReservation" : 0.0,
        "NumCpu" : 2,
        "CPUCycleEntlMin" : 0.0,
        "RecommendedMemReservation" : 0.0,
        "RecommendedMem" : 2.0,
        "MemCapacity" : 4.0
    },
    "MemSize" : "oversized",
    "CPUSize" : "oversized"
} ]
}
```

To get sizing information of all the datastores that are part of a specific cluster

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/cluster/<instance_
id>/datastoresizing
```

The following JSON is returned in the response:

```
{
    "instancecount" : 5,
    "instances" : [
```

```
{
    "properties" : {
        "ParentUUID" : "35353636-3235-4753-4834-303241525843",
        "SystemVirtType" : "Hyper-V",
        "SystemID" : "02e045f8-3ea8-4d7d-844b-eb82ea0550bb",
        "ParentType" : "HOST",
        "ClusterName" : "",
        "BelongsToDatacenter" : "",
        "Type" : "NTFS",
        "SystemRole" : "DATASTORE",
        "LSName" : "",
        "SystemName" : "New Volume (E:)"
    },
    "resourcepath" : "/datastore/1",
    "resourceid" : 1,
    "sizing" :
    {
        "PO_StorageReclaimable" : 0.0,
        "PO_StorageAllocationPossibleAfterReclaim" : 48130.54,
        "PO_DiskUsage" : 917.57
    }
}, .....
```

Host

The following section gives sample JSON response generated by the REST API for optimization recommendations for Host.

To get information of all the Hosts

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/host
```

The following JSON is was returned in the response:

```
{
  "instancecount" : 8,
  "instances" : [
    {
      "properties" :
      {
        "SystemOSName" : "Microsoft Windows Server 2012 R2 Datacenter",
        "StatTime" : "",
        "SystemID" : "35353636-3235-4753-4834-303241525843",
        "SystemOSRelease" : "",
        "ParentType" : "CLUSTER",
        "SystemPath" : "",
        "SystemRole" : "HOST",
        "Type" : "",
        "UseHardwareAssistedVirtualization" : "",
        "MacAddress" : "2C:44:FD:87:C2:30",
        "SystemHostHostName" : "",
        "SystemVirtType" : "Hyper-V",
        "ClusterName" : "hyperv-clus",
        "SystemOSVersion" : "6.3.9600",
        "ManagementIP" : "",
        "FlavorId" : "",
        "VCIPAddress" : "",
        "SystemVirtPlatform" : "",
        "InstanceName" : "",
        "CPUArch" : "",
```

```
        "ProcessorVendor" : "GenuineIntel",
        "VirtVersion" : "6.3.9600",
        "SystemName" : "GOLDRICH6",
        "MarkedForDeletion" : "",
        "VMOwnerId" : "",
        "HPCSPort" : "",
        "LSID" : "",
        "HAEnabled" : "",
        "ParentUUID" : "hyperv-clus",
        "DynamicMemoryEnabled" : "",
        "LSMode" : "",
        "CPUVendor" : "",
        "HasSnapshot" : "",
        "BelongsToDatacenter" : "",
        "VMHost_UUID" : "",
        "LSName" : "",
        "VirtVersionState" : "",
        "ManagementUUID" : "",
        "LSShared" : "",
        "HPCSProtocol" : "",
        "SystemHostName" : "GOLDRICH6.iwflabs.com"
    },
    "resourcepath" : "/host/1",
    "resourceid" : 1
},....
```

To get information of a specific Host

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/host/<instance_id>
```

The following JSON is returned in the response:

```
{
  "instances" : [
    {
      "properties" :
      {
        "SystemOSName" : "VMware ESXi",
        "StatTime" : "",
        "SystemID" : "31353337-3135-4753-4834-32344b315944",
        "SystemOSRelease" : "",
        "ParentType" : "CLUSTER",
        "SystemPath" : "",
        "SystemRole" : "HOST",
        "Type" : "",
        "UseHardwareAssistedVirtualization" : "",
        "MacAddress" :
        "fc:15:b4:22:70:90;fc:15:b4:22:70:91;fc:15:b4:22:70:92;fc:15:b4:22:70:94;fc:15:b4:22:70:95;fc:15:b4:22:70:96;",
        "SystemHostHostName" : "",
        "SystemVirtType" : "VMWARE",
        "ClusterName" : "BLR-APPLE-CLUSTER",
        "SystemOSVersion" : "",
        "ManagementIP" : "",
        "FlavorId" : "",
        "VCIPAddress" : "16.184.45.37",
        "SystemVirtPlatform" : "",
        "InstanceName" : "",
        "CPUArch" : "",
```

```
        "ProcessorVendor" : "",
        "VirtVersion" : "",
        "SystemName" : "16.184.42.74",
        "MarkedForDeletion" : "",
        "VMOwnerId" : "",
        "HPCSPort" : "",
        "LSID" : "39",
        "HAEnabled" : "",
        "ParentUUID" : "16.184.45.37/BLR-APPLE-DC/BLR-APPLE-CLUSTER",
        "DynamicMemoryEnabled" : "",
        "LSMode" : "Capped",
        "CPUVendor" : "",
        "HasSnapshot" : "",
        "BelongsToDatacenter" : "16.184.45.37/BLR-APPLE-DC",
        "VMHost_UUID" : "",
        "LSName" : "31353337-3135-4753-4834-32344b315944",
        "VirtVersionState" : "",
        "ManagementUUID" : "",
        "LSShared" : "Dedicated",
        "HPCSProtocol" : "",
        "SystemHostName" : "16.184.42.74"
    },
    "resourcepath" : "/host/63",
    "resourceid" : 63
} ]
}
```


To get sizing information of all the VMs in a specific Host

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/host/<instance_id>/vmsizing
```

The following JSON is returned in the response:

```
{
  "instancecount" : 2,
  "instances" : [
    {
      "properties" :
      {
        "SystemOSName" : "CentOS 4/5/6 (64-bit)",
        "StatTime" : "",
        "SystemID" : "421a3f0d-4bfe-0808-d643-09c48f9a9bed",
        "SystemOSRelease" : "",
        "ParentType" : "RESPOOL",
        "SystemPath" : "[datastore1]Paul_33_210_18092014/Paul_33_210_18092014.vmx",
        "SystemRole" : "GUEST",
        "Type" : "",
        "UseHardwareAssistedVirtualization" : "",
        "MacAddress" : "00:50:56:9a:5a:77,52:54:00:a6:13:3d",
        "SystemHostHostName" : "16.184.42.74",
        "SystemVirtType" : "VMWARE",
        "ClusterName" : "BLR-APPLE-CLUSTER",
        "SystemOSVersion" : "",
        "ManagementIP" : "",
        "FlavorId" : ""
      }
    }
  ]
}
```

```
    "VCIPAddress" : "",
    "SystemVirtPlatform" : "",
    "InstanceName" : "",
    "CPUArch" : "",
    "ProcessorVendor" : "",
    "VirtVersion" : "",
    "SystemName" : "Paul_33_210_18092014",
    "MarkedForDeletion" : "",
    "VMOwnerId" : "",
    "HPCSPort" : "",
    "LSID" : "71",
    "HAEnabled" : "",
    "ParentUUID" : "16.184.45.37/BLR-APPLE-DC/BLR-APPLE-CLUSTER/Apri",
    "DynamicMemoryEnabled" : "",
    "LSMode" : "UnCapped",
    "CPUVendor" : "",
    "HasSnapshot" : "",
    "BelongsToDatacenter" : "16.184.45.37/BLR-APPLE-DC",
    "VMHost_UUID" : "",
    "LSName" : "421a3f0d-4bfe-0808-d643-09c48f9a9bed",
    "VirtVersionState" : "",
    "ManagementUUID" : "",
    "LSShared" : "Shared",
    "HPCSProtocol" : "",
    "SystemHostName" : "martellvm26.ind.hp.com"
  },
  "resourcepath" : "/vm/70",
  "resourceid" : 70,
```

```
"sizing" : {  
    "PO_MemCapacity" : 16.0,  
    "PO_CPUCycleEntlMin" : 0.0,  
    "PO_RecommendedNumCPU" : 2.0,  
    "PO_RecommendedCPUReservation" : 0.0,  
    "PO_RecommendedMem" : 6.0,  
    "PO_MemEntlMin" : 0.0,  
    "PO_RecommendedMemReservation" : 0.0,  
    "PO_NumCpu" : 4  
},  
"MemSize" : "oversized",  
"CPUSize" : "oversized"  
},....
```

To get sizing information of a specific VM in a Host filtered based on name and sizing parameters

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/host/<instance_id>/vmsizing?name=<system_name>&sizing=<sizing>
```

```
{  
"instancecount" : 1,  
"instances" : [ {  
    "properties" : {  
        "SystemOSName" : "CentOS 4/5/6 (64-bit)",  
        "StatTime" : "",  
        "SystemID" : "420ccb5b-2e36-7fa2-13a5-e39980e/goldrich8vm5.ind.hp.com_vPV2.10.005_4th_NOV",  
        "SystemPath" : "[Datastore2]goldrich8vm5.ind.hp.com_vPV2.10.005_4th_NOV/goldrich8vm5.ind.hp.com_vPV2.10.005_4th_NOV.vmx",
```

```
"SystemOSRelease" : "",
"ParentType" : "HOST",
"Type" : "",
"SystemRole" : "GUEST",
"UseHardwareAssistedVirtualization" : "",
"SystemHostHostName" : "ovpesx10.ind.hp.com",
"MacAddress" : "00:50:56:8c:59:02,52:54:00:45:76:34",
"SystemVirtType" : "VMWARE",
"ClusterName" : "ispi_qa_CLUSTER_1029",
"SystemOSVersion" : "",
"ManagementIP" : "16.184.40.94",
"FlavorId" : "",
"VCIPAddress" : "",
"SystemVirtPlatform" : "",
"CPUArch" : "",
"InstanceName" : "",
"ProcessorVendor" : "",
"VirtVersion" : "",
"SystemName" : "goldrich8vm5.ind.hp.com_vPV2.10.005_4th_NOV",
"MarkedForDeletion" : "",
"VMOwnerId" : "",
"HPCSPort" : "381",
"LSID" : "1442",
"HAEnabled" : "",
"ParentUUID" : "34353139-3933-5347-4838-343658434531",
"DynamicMemoryEnabled" : "",
"LSMode" : "UnCapped",
"CPUVendor" : "",
```

```
        "HasSnapshot" : "",
        "BelongsToDatacenter" : "iwfvm01029.hpswlab.s.adapps.hp.com/ispi_qa_DC_
1029",
        "VMHost_UUID" : "",
        "LSName" : "420ccb5b-2e36-7fa2-13a5-e39980e66a85/goldrich8vm5.ind.hp.com_
vPV2.10.005_4th_NOV",
        "VirtVersionState" : "",
        "LSShared" : "Shared",
        "ManagementUUID" : "40120627-5B31-4004-BF28-0050568C5902",
        "SystemHostName" : "goldrich8vm5.ind.hp.com",
        "HPCSProtocol" : "http"
    },
    "resourcepath" : "/vm/5",
    "resourceid" : 5,
    "sizing" : {
        "MemEntlMin" : 0.0,
        "RecommendedNumCPU" : 2.0,
        "RecommendedCPUReservation" : 0.0,
        "NumCpu" : 4,
        "CPUCycleEntlMin" : 0.0,
        "RecommendedMemReservation" : 0.0,
        "RecommendedMem" : 6.0,
        "MemCapacity" : 16.0
    },
    "MemSize" : "oversized",
    "CPUSize" : "oversized"
} ]
}
```

To get datastore information and sizing details of all the datastores in a specific Host

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/<instance_id>/datastoresizing
```

The following JSON is returned in the response:

```
{
  "instance" : [
    {
      "properties" :
      {
        "ParentUUID" : "16.184.45.37/SMV-BLADE",
        "SystemVirtType" : "VMWARE",
        "IsFullyCached" : "",
        "SystemID" : "16.184.45.37/SMV-BLADE/SMV-Cluster",
        "BelongsToDatacenter" : "16.184.45.37/SMV-BLADE",
        "ParentType" : "DATACENTER",
        "Type" : "VMWare DRS",
        "SystemRole" : "CLUSTER",
        "LSName" : "16.184.45.37/SMV-BLADE/SMV-Cluster",
        "SystemName" : "SMV-Cluster",
        "MarkedForDeletion" : "",
        "SystemVirtPlatform" : ""
      },
      "resourcepath" : "/cluster/2",
      "resourceid" : 2
    } ]
}
```

Virtual Machines

The following section gives sample JSON response generated by the REST API for optimization recommendations for Virtual Machines.

To get information about all the VMs along with sizing details

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/vm
```

The following JSON is returned in the response:

```
{
  "instancecount" : 107,
  "instances" : [
    {
      "properties" : {
        "SystemOSName" : "Red Hat Enterprise Linux 6 (64-bit)",
        "StatTime" : "",
        "SystemID" : "421ad522-d7f5-0025-1b0a-409df1a2898f",
        "SystemOSRelease" : "",
        "ParentType" : "HOST",
        "SystemPath" : "[BLR-LUN-0112]
martellvm40158.hpswlab.com/martellvm40158.hpswlab.com.vmx",
        "SystemRole" : "GUEST",
        "Type" : "",
        "UseHardwareAssistedVirtualization" : "",
        "MacAddress" : "00:50:56:9a:79:aa",
        "SystemHostHostName" : "agentesx.ind.hp.com",
        "SystemVirtType" : "VMWARE",
        "ClusterName" : "BLR-APPLE-CLUSTER",
        "SystemOSVersion" : "",
```

```
"ManagementIP" : "",  
"FlavorId" : "",  
"VCIPAddress" : "",  
"SystemVirtPlatform" : "",  
"InstanceName" : "",  
"CPUArch" : "",  
"ProcessorVendor" : "",  
"VirtVersion" : "",  
"SystemName" : "martellvm40158.hpswlab.hp.com",  
"MarkedForDeletion" : "",  
"VMOwnerId" : "",  
"HPCSPort" : "",  
"LSID" : "79",  
"HAEnabled" : "",  
"ParentUUID" : "38393636-3430-4753-4832-333442364c53",  
"DynamicMemoryEnabled" : "",  
"LSMode" : "UnCapped",  
"CPUVendor" : "",  
"HasSnapshot" : "",  
"BelongsToDatacenter" : "16.184.45.37/BLR-APPLE-DC",  
"VMHost_UUID" : "",  
"LSName" : "421ad522-d7f5-0025-1b0a-409df1a2898f",  
"VirtVersionState" : "",  
"ManagementUUID" : "",  
"LSShared" : "Shared",  
"HPCSProtocol" : "",  
"SystemHostName" : "ci-0050569A79AA"  
},
```



```
"resourcepath" : "/vm/62",
"resourceid" : 62,
"sizing" : {
    "PO_MemCapacity" : 10.0,
    "PO_CPUCycleEntlMin" : 0.0,
    "PO_RecommendedNumCPU" : 1.0,
    "PO_RecommendedCPUReservation" : 0.0,
    "PO_RecommendedMem" : 5.0,
    "PO_MemEntlMin" : 0.0,
    "PO_RecommendedMemReservation" : 0.0,
    "PO_NumCpu" : 2
},
"MemSize" : "oversized",
"CPUSize" : "oversized"
},
.....
```

To get information of a specific VM

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/<instance_id>
```

The following JSON is returned in the response:

```
{
    "instance" : [
        {
            "properties" : {
                "SystemOSName" : "Windows Server 2012 R2 Standard",
                "StatTime" : "",
                "SystemID" : "1FEDE220-B72F-432E-AFA2-6FCE94D19B1E",
```

```
"SystemOSRelease" : "",
"ParentType" : "HOST",
"SystemPath" : "",
"SystemRole" : "GUEST",
"Type" : "",
"UseHardwareAssistedVirtualization" : "",
"MacAddress" : "00:15:5D:2F:78:13",
"SystemHostHostName" : "GOLDRICH6.iwflabs.com",
"SystemVirtType" : "Hyper-V",
"ClusterName" : "hyperv-clus",
"SystemOSVersion" : "",
"ManagementIP" : "",
"FlavorId" : "",
"VCIPAddress" : "",
"SystemVirtPlatform" : "",
"InstanceName" : "",
"CPUArch" : "",
"ProcessorVendor" : "",
"VirtVersion" : "",
"SystemName" : "HyperV_VM2",
"MarkedForDeletion" : "",
"VMOwnerId" : "",
"HPCSPort" : "",
"LSID" : "",
"HAEnabled" : "",
"ParentUUID" : "35353636-3235-4753-4834-303241525843",
"DynamicMemoryEnabled" : "True",
"LSMode" : "",
```

```
        "CPUVendor" : "",
        "HasSnapshot" : "",
        "BelongsToDatacenter" : "",
        "VMHost_UUID" : "",
        "LSName" : "",
        "VirtVersionState" : "",
        "ManagementUUID" : "",
        "LSShared" : "",
        "HPCSProtocol" : "",
        "SystemHostName" : "WIN-A80SLN3KECF"
    },
    "resourcepath" : "/vm/5",
    "resourceid" : 5,
    "sizing" : {
        "PO_MemCapacity" : 1024.0,
        "PO_CPUCycleEntlMin" : 0.0,
        "PO_RecommendedNumCPU" : 0.0,
        "PO_RecommendedCPUReservation" : 0.0,
        "PO_RecommendedMem" : 206.0,
        "PO_MemEntlMin" : 0.5,
        "PO_RecommendedMemReservation" : 0.548828125,
        "PO_NumCpu" : 1
    },
    "MemSize" : "oversized",
    "CPUSize" : "oversized"
} ]
}
}
```

Datastores

The following section gives sample JSON response generated by the REST API for optimization recommendations for Datastores.

To get datastore information along with sizing details

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datastore
```

The following JSON is returned in the response:

```
{
  "instancecount" : 20,
  "instances" : [ {
    "properties" :
    {
      "ParentUUID" : "35353636-3235-4753-4834-303241525843",
      "SystemVirtType" : "Hyper-V",
      "SystemID" : "02e045f8-3ea8-4d7d-844b-eb82ea0550bb",
      "ParentType" : "HOST",
      "ClusterName" : "",
      "BelongsToDatacenter" : "",
      "Type" : "NTFS",
      "SystemRole" : "DATASTORE",
      "LSName" : "",
      "SystemName" : "New Volume (E:)"
    },
    "resourcepath" : "/datastore/1",
    "resourceid" : 1,
    "sizing" :
    {
```

```
        "PO_StorageReclaimable" : 0.0,
        "PO_StorageAllocationPossibleAfterReclaim" : 48130.54,
        "PO_DiskUsage" : 917.57
    }
}, .....
{
    "properties" :
    {
        "ParentUUID" : "35353636-3235-4753-4834-303241525843",
        "SystemVirtType" : "Hyper-V",
        "SystemID" : "1dfeb6df-4da3-11e3-80b4-806e6f6e6963",
        "ParentType" : "HOST",
        "ClusterName" : "",
        "BelongsToDatacenter" : "",
        "Type" : "NTFS",
        "SystemRole" : "DATASTORE",
        "LSName" : "",
        "SystemName" : "GOLDRICH6 (C:)"
    },
    "resourcepath" : "/datastore/2",
    "resourceid" : 2,
    "sizing" : {
        "PO_StorageReclaimable" : 0.0,
        "PO_StorageAllocationPossibleAfterReclaim" : 4767.22,
        "PO_DiskUsage" : 86.3
    }
}
```

To get information of a specific datastore

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datastore/<instance_id>
```

The following JSON is returned in the response:

```
{
  "instance" : [
    {
      "properties" :
      {
        "ParentUUID" : "35353636-3235-4753-4834-303241525843",
        "SystemVirtType" : "Hyper-V",
        "SystemID" : "1dfeb6df-4da3-11e3-80b4-806e6f6e6963",
        "ParentType" : "HOST",
        "ClusterName" : "",
        "BelongsToDatacenter" : "",
        "Type" : "NTFS",
        "SystemRole" : "DATASTORE",
        "LSName" : "",
        "SystemName" : "GOLDRICH6 (C:)"
      },
      "resourcepath" : "/datastore/2",
      "resourceid" : 2,
      "sizing" :
      {
        "PO_StorageReclaimable" : 0.0,
        "PO_StorageAllocationPossibleAfterReclaim" : 4767.22,
```

```
        "PO_DiskUsage" : 86.3
    }
} ]
}
```

To get information about all the VMs on a specific datastore

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datastore/<instance_
id>/vm
```

The following JSON is returned in the response:

```
{
  "instancecount" : 3,
  "instances" : [ {
    "properties" : {
      "ParentUUID" : "4ba37e78-d1000676-135b-d8d385ab1163",
      "SystemVirtType" : "VMWARE",
      "SystemID" : "42131097-1760-a71e-87b6-8b3a60178fb4/JMeter4vPVTest",
      "ParentType" : "DATASTORE",
      "SystemRole" : "GUEST",
      "SystemName" : "JMeter4vPVTest",
      "DatastoreID" : "4ba37e78-d1000676-135b-d8d385ab1163",
      "DatastoreName" : "EVA_HPSW-Template002"
    },
    "resourcepath" : "/vm/74",
    "resourceid" : 74,
    "sizing" : {
      "DiskProvisioned" : 18.1,
      "ActivityStatus" : "idle",
```

```
        "DiskUsed" : 4.55
    },
    "datastorepath" : ""
},
{
    "properties" : {
        "ParentUUID" : "4ba37e78-d1000676-135b-d8d385ab1163",
        "SystemVirtType" : "VMWARE",
        "SystemID" : "421339d3-4325-e674-a845-f2d38c885045/pmivpv-117-2.1-5-NOV",
        "ParentType" : "DATASTORE",
        "SystemRole" : "GUEST",
        "SystemName" : "pmivpv-117-2.1-5-NOV",
        "DatastoreID" : "4ba37e78-d1000676-135b-d8d385ab1163",
        "DatastoreName" : "EVA_HPSW-Template002"
    },
    "resourcepath" : "/vm/194",
    "resourceid" : 194,
    "sizing" : {
        "DiskProvisioned" : 56.09,
        "ActivityStatus" : "active",
        "DiskUsed" : 22.4
    },
    "datastorepath" : ""
}, {
    "properties" : {
        "ParentUUID" : "4ba37e78-d1000676-135b-d8d385ab1163",
        "SystemVirtType" : "VMWARE",
        "SystemID" : "42139d3b-6e0b-7fbd-e55e-b6580e06ff8a/vpv-kamboji-2.1-13-Nov"
```



```
        "ParentType" : "DATASTORE",
        "SystemRole" : "GUEST",
        "SystemName" : "vpv-kamboji-2.1-13-Nov",
        "DatastoreID" : "4ba37e78-d1000676-135b-d8d385ab1163",
        "DatastoreName" : "EVA_HPSW-Template002"
    },
    "resourcepath" : "/vm/509",
    "resourceid" : 509,
    "sizing" : { },
    "datastorepath" : ""
} ]
}
```

To get information of a specific VM in a datastore filtered on the basis of Status

Use the following URL:

```
https://<ip_address>:<port_number>/PV/api/v1/datastore/<instance_id>/vm?status=<vm_status>
```

The following JSON is returned in the response:

```
{
  "instancecount" : 1,
  "instances" : [ {
    "properties" : {
      "ParentUUID" : "4ba37e78-d1000676-135b-d8d385ab1163",
      "SystemVirtType" : "VMWARE",
      "SystemID" : "42131097-1760-a71e-87b6-8b3a60178fb4/JMeter4vPVTest",
      "ParentType" : "DATASTORE",
      "SystemRole" : "GUEST",
      "SystemName" : "JMeter4vPVTest",
```

```
        "DatastoreID" : "4ba37e78-d1000676-135b-d8d385ab1163",  
        "DatastoreName" : "EVA_HPSW-Template002"  
    },  
    "resourcepath" : "/vm/74",  
    "resourceid" : 74,  
    "sizing" : {  
        "DiskProvisioned" : 18.1,  
        "ActivityStatus" : "idle",  
        "DiskUsed" : 4.55  
    },  
    "datastorepath" : ""  
} ]  
}
```

Using REST API to view Powered Off VMs

HPE Cloud Optimizer provides REST APIs to view the details of powered off VMs in your virtualized environment.

Resource Model

The REST API to view powered off VMs can be invoked at the following levels:

- All domains
- vCenter
- Datacenter
- Host
- Datacenter > Cluster
- Datacenter > Host
- Datacenter > Cluster > Host

Resource Information

Response formats	JSON
------------------	------

Request Header requirements	Content-Type: application/json Authorization: Basic <Base 64 encoded username:password> (optional, required only when LDAP is enabled)
Supported request methods	POST

Resource URL

The Base URL to get information about powered off VMs is:

`https://<ip_address or host_name>:<port_number>/PV/api/v1`

In this instance,

- `ip_address` or `host_name` is the IP address or host name of the HPE Cloud Optimizer server.
- `port_number` is the port number of the HPE Cloud Optimizer server. The default port number is 8444.

The Base URL must be appended with appropriate URIs to get information about the powered off VMs in the environment. The following table summarizes REST API URL details:

API	URL
To view the details of all the powered off VMs in the virtualized environment.	<code>https://<ip_address>:<port_number>/PV/api/v1/poweredoffvmlist</code>
To view the details of powered off VMs in a specific Datacenter.	<code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/poweredoffvmlist</code> In this instance, <code>datacenter_id</code> is the ID of the datacenter.
To view the details of powered off VMs in a specific host in a Datacenter.	<code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/host/<host_id>/poweredoffvmlist</code> In this instance,

API	URL
	<p>datacenter_id is the ID of the datacenter.</p> <p>host_id is the ID of the host.</p>
To view the details of powered off VMs in a specific cluster in a Datacenter.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/cluster/<cluster_id>/host/<host_id>/poweredoffvmlist</code></p> <p>In this instance,</p> <p>datacenter_id is the ID of the datacenter.</p> <p>host_id is the ID of the host.</p> <p>cluster_id is the ID of the Cluster.</p>
To view the details of powered off VMs in a specific host that belongs to cluster in a Datacenter.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/cluster/<cluster_id>poweredoffvmlist</code></p> <p>In this instance,</p> <p>datacenter_id is the ID for that datacenter.</p> <p>cluster_id is the ID of the Cluster.</p>
To view the details of powered off VMs in a specific host.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/host/<host_id>/poweredoffvmlist</code></p> <p>In this instance,</p> <p>host_id is the ID of the host.</p>
To view the details of powered off VMs in all the vCenters in the environment.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/vcenters/poweredoffvmlist</code></p>

JSON Response

The following JSON is returned in the response:

```
{
```

```
"poweredoffcount" : n,
"poweredoffList" :
[ {
  "systemName" : "vPV_361_16.184.40.70",
  "annotations" : "",
  "poweredOffSince" : 1430980102,
  "downTimePercent" : 100.0,
  "hostName" : "16.184.42.71",
  "diskUsed" : 21092.0,
  "cpuutil" : 0.0,
  "memutil" : 0.0,
  "diskIO" : 0.0,
  "netByte" : 0.0,
},
{
  "systemName" : "RHEL-KVMHost-vm2",
  "annotations" : "",
  "poweredOffSince" : 1430980102,
  "downTimePercent" : 100.0,
  "hostName" : "16.184.42.71",
  "diskUsed" : 230400.0,
  "cpuutil" : 0.0,
  "memutil" : 0.0,
  "diskIO" : 0.0,
  "netByte" : 0.0,
}
.
```

```
.  
{  
    "systemName" : "RHEL-KVMHost-vm3",  
    "annotations" : "",  
    "poweredOffSince" : 1430980102,  
    "downTimePercent" : 100.0,  
    "hostName" : "16.184.42.71",  
    "diskUsed" : 230400.0,  
    "cpuutil" : 0.0,  
    "memutil" : 0.0,  
    "diskIO" : 0.0,  
    "netByte" : 0.0,  
}  
} ]
```

Using REST APIs to view Idle VMs

HPE Cloud Optimizer provides REST APIs to view the details of Idle VMs in your virtualized environment.

Resource Model

The REST API to view idle VMs can be invoked at the following levels:

- All domains
- vCenter
- Datacenter
- Host
- Datacenter > Cluster
- Datacenter > Host
- Datacenter > Cluster > Host

Resource Information

Response formats	JSON
Request Header requirements	Content-Type: application/json Authorization: Basic <Base 64 encoded username:password> (optional, required only when LDAP is enabled)
Supported request methods	POST

Resource URL

The Base URL to get information about Idle VMs is:

`https://<ip_address or host_name>:<port_number>/PV/api/v1`

In this instance,

- `ip_address` or `host_name` is the IP address or host name of the HPE Cloud Optimizer server.
- `port_number` is the port number of the HPE Cloud Optimizer server. The default port number is 8444.

The Base URL must be appended with appropriate URIs to get information about the Idle VMs in the environment. The following table summarizes REST API URL details:

API	URL
To view the details of Idle VMs in the virtualized environment.	<code>https://<ip_address>:<port_number>/PV/api/v1/idlevmlist</code>
To view the details of Idle VMs in a specific Datacenter.	<code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/idlevmlist</code> In this instance, <code>datacenter_id</code> is the ID of the datacenter.
To view the details of Idle VMs in a specific host in a Datacenter.	<code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/host/<host_id>/idlevmlist</code> In this instance,

API	URL
	<p>datacenter_id is the ID of the datacenter.</p> <p>host_id is the ID of the host.</p>
To view the details of Idle VMs in a specific cluster in a Datacenter.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/cluster/<cluster_id>/host/<host_id>/idlevm</code></p> <p>In this instance,</p> <p>datacenter_id is the ID for that datacenter.</p> <p>host_id is the ID of the host.</p> <p>cluster_id is the ID of the Cluster.</p>
To view the details of Idle VMs in a specific host that belongs to cluster in a Datacenter.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/cluster/<cluster_id>idlevm</code></p> <p>In this instance,</p> <p>datacenter_id is the ID for that datacenter.</p> <p>cluster_id is the ID of the Cluster.</p>
To view the details of Idle VMs in a specific host.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/host/<host_id>/idlevm</code></p> <p>In this instance,</p> <p>datacenter_id is the ID for that datacenter.</p> <p>host_id is the ID of the host.</p> <p>cluster_id is the ID of the Cluster.</p>
To view the details of Idle VMs in all the vCenters in the environment.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/vcenters/idlevm</code></p>

JSON Response

The following JSON is returned in the response:


```
{
  "idlecount" : n,
  "idleList" :
  [ {
    "systemName" : "vccn_do_not_delete-161",
    "annotations" : "",
    "poweredOffSince" : 0,
    "downTimePercent" : 0.0,
    "hostName" : "",
    "diskUsed" : 28602.0,
    "cpuutil" : 0.17,
    "memutil" : 1.44,
    "diskIO" : 8.0,
    "netByte" : 2.0,
  },
  {
    "systemName" : "smvblade5vm9.hpswlab.hp.com",
    "annotations" : "",
    "poweredOffSince" : 0,
    "downTimePercent" : 0.0,
    "hostName" : "",
    "diskUsed" : 78391.0,
    "cpuutil" : 0.0,
    "memutil" : 0.0,
    "diskIO" : 0.0,
    "netByte" : 0.0,
  }
}
```

```
.  
.br/>{  
  "systemName" : "16.184.45.134",  
  "annotations" : "",  
  "poweredOffSince" : 0,  
  "downTimePercent" : 0.0,  
  "hostName" : "",  
  "diskUsed" : 11356.0,  
  "cpuutil" : 0.01,  
  "memutil" : 4.99,  
  "diskIO" : 1.0,  
  "netByte" : 0.0,  
}  
} ]
```

Using REST APIs to view CPU, Memory, and Storage allocation of VMs

HPE Cloud Optimizer provides REST APIs to get the CPU, Memory, and Storage allocation of a set of VMs over a specific time interval. You can use this information to calculate the billing cost for the CPU, Memory, and Storage resources used by your business applications and services.

Resource Information

Response formats	JSON
Request Header requirements	Content-Type: application/json Authorization: Basic <Base 64 encoded username:password> (optional, required only when LDAP is enabled)

Supported request methods	GET
---------------------------	-----

Resource URL

The Base URL to get information about VMs is:

`https://<IP_Address or hostname>:<port_number>/PV/api/v1`

In this instance,

- `IP_Address` or `hostname` is the IP address or host name of the HPE Cloud Optimizer server.
- `port_number` is the port number of the HPE Cloud Optimizer server. The default port number is 8444.

The Base URL must be appended with appropriate URIs to get the allocation information of VMs. The following table summarizes REST API URL details at various levels:

API	URL
To view the allocation details of all the VMs in the environment.	<p><code>https://<IP_Address or hostname>:<port_number>/PV/api/v1/vm/utilization?starttime=<starttime>&endtime=<endtime>&start=<start-index>&limit=<limit></code></p> <p>In this instance,</p> <p><code><starttime></code> is the start time from which you want to view the VM information. The <code>starttime</code> must be specified as an Epoch timestamp.</p> <p><code><endtime></code> is the time till which you want to view the VM information. The <code>endtime</code> must be specified as an Epoch timestamp.</p> <p><code><start-index></code> is the starting number for the VM instance count.</p> <p><code><limit></code> is the limiting number for the VM instance count.</p> <p>All the parameters are mandatory. Click here to view the list of metrics in this API.</p>
To view the allocation details of a specific VM	<p><code>https://<IP_Address or hostname>:<port_number>/PV/api/v1/vm/<Resource-ID>/utilization?starttime=<starttime>&endtime=<endtime></code></p> <p>In this instance,</p> <p><code><Resource-ID></code> is the instance ID of the VM.</p>

API	URL
	<p><starttime> is the start time from which you want to view the VM information. The starttime must be specified as an Epoch timestamp.</p> <p><endtime> is the time till which you want to view the VM information. The endtime must be specified as an Epoch timestamp.</p> <p>All the parameters are mandatory.</p>
To view the allocation details of all the Business Groups	<p>https://<vpv_hostname>:8444/PV/api/v1/group/utilization?starttime=<starttime>&endtime=<endtime></p>
To view the allocation details of VMs in a specific Business Group.	<p>https://<IP_Address or hostname>:<port_number>/PV/api/v1/vm/group/<grouping_name>/<Resource-ID>/utilization?starttime=<starttime>&endtime=<endtime>&start=<start-index>&limit=<limit></p> <p>In this instance,</p> <p><grouping_name> is name of the Business Grouping to which the Business Group belongs.</p> <p><Resource-ID> is the instance ID of the Business Group.</p> <p><starttime> is the start time from which you want to view the VM information. The starttime must be specified as an Epoch timestamp.</p> <p><endtime> is the time till which you want to view the VM information. The endtime must be specified as an Epoch timestamp.</p> <p><start-index> is the starting number for the VM instance count.</p> <p><end-index> is the limiting number for the VM instance count.</p> <p>All the parameters are mandatory.</p>
To view the allocation details of VMs	<p>https://<IP_Address or hostname>:<port_number>/PV/api/v1/host/<Resource-ID>/vm/utilization?starttime=<starttime>&endtime=<endtime>&start=<start-index>&limit=<limit></p>

API	URL
in a specific Host.	<p>In this instance,</p> <p><Resource-ID> is the instance ID of the Host.</p> <p><starttime> is the time from which you want to view the VM information. The starttime must be specified as an Epoch timestamp.</p> <p><endtime> is the time till which you want to view the VM information. The endtime must be specified as an Epoch timestamp.</p> <p><start-index> is the starting number for the VM instance count.</p> <p><end-index> is the limiting number for the VM instance count.</p> <p>All the parameters are mandatory.</p>
To view the allocation details of VMs in a specific cluster.	<p><a href="https://<IP_Address or hostname>:<port_number>/PV/api/v1/cluster/<Resource-ID>/vm/utilization?starttime=<starttime>&endtime=<endtime>&start=<start-index>&limit=<limit>">https://<IP_Address or hostname>:<port_number>/PV/api/v1/cluster/<Resource-ID>/vm/utilization?starttime=<starttime>&endtime=<endtime>&start=<start-index>&limit=<limit></p> <p>In this instance,</p> <p><Resource-ID> is the instance ID of the cluster.</p> <p><starttime> is the start time from which you want to view the VM information. The starttime must be specified as an Epoch timestamp.</p> <p><endtime> is the time till which you want to view the VM information. The endtime must be specified as an Epoch timestamp.</p> <p><start-index> is the starting number for the VM instance count.</p> <p><end-index> is the limiting number for the VM instance count</p> <p>All the parameters are mandatory.</p>
To view the allocation details of VMs in a specific Datacenter.	<p><a href="https://<IP_Address or hostname>:<port_number>/PV/api/v1/datacenter/<Resource-ID>/vm/utilization?starttime=<time>&endtime=<time>&start=<start-index>&limit=<limit>">https://<IP_Address or hostname>:<port_number>/PV/api/v1/datacenter/<Resource-ID>/vm/utilization?starttime=<time>&endtime=<time>&start=<start-index>&limit=<limit></p> <p>In this instance,</p> <p><Resource-ID> is the instance ID of the Datacenter.</p> <p><starttime> is the start time from which you want to view the VM</p>

API	URL
	<p>information. The <code>starttime</code> must be specified as an Epoch timestamp.</p> <p><code><endtime></code> is the time till which you want to view the VM information. The <code>endtime</code> must be specified as an Epoch timestamp.</p> <p><code><start-index></code> is the starting number for the VM instance count.</p> <p><code><end-index></code> is the limiting number for the VM instance count</p> <p>All the parameters are mandatory.</p>

Note: In case of data gap in the selected time interval, the utilization is, by default, considered as zero for the time interval that has no data. You can configure HPE Cloud Optimizer, to consider the earlier utilization for that time interval.

To configure HPE Cloud Optimizer to consider the earlier utilization if there is a data gap in the selected time interval, perform the following steps:

1. On the HPE Cloud Optimizer machine, log on as root.
2. Run the following command:

```
<install_dir>/bin/ovconfchg -edit
```

In this instance, `install_dir` is the directory where HPE Cloud Optimizer is installed. The default location is `/opt/OV`.

A text file opens.
3. In the file, type the following to create a namespace:

```
[pvcd.rest]
```
4. In the `[pvcd.rest]` namespace, type the following:

```
DataGapValue=before
```

This ensures that the earlier utilization is considered, if there is a data gap in the selected time interval.
5. Save and close the file.

Note: All the values in the text file are case sensitive.

Example

For example, there are 350 VM instances in the environment. You want to view 100 VM instances in one REST API call.

To achieve this, use the following REST API calls in sequence:

1. `https://xx.xxx.xx.xxx:8444/PV/api/v1/vm/utilization?starttime=1441305000&endtime=1441391400&start=1&limit=100`

2. `https://xx.xxx.xx.xxx:8444/PV/api/v1/vm/utilization?starttime=1441305000&endtime=1441391400&start=101&limit=100`
3. `https://xx.xxx.xx.xxx:8444/PV/api/v1/vm/utilization?starttime=1441305000&endtime=1441391400&start=201&limit=100`
4. `https://xx.xxx.xx.xxx:8444/PV/api/v1/vm/utilization?starttime=1441305000&endtime=1441391400&start=301&limit=100`

As the total number of instances are 350, the last REST API call would return only 49 instances. Hence, you can conclude that there are no more VM instances left and no more API calls are required.

JSON Response

The following JSON is returned in the response:

For all the VMs in the environment

```
{
  instancecount: 100
  instances: [100]
  0: {
    properties: {
      SystemVirtType: "VMWARE"
      SystemID: "421a192e-fde2-92ae-90b7-ba4e976a7fd1"
      SystemName: "RHEL-KVMHost-1"
      SystemOSName: "Red Hat Enterprise Linux 6 (64-bit)"
      SystemHostName: "vpvqa8.ind.sample.com"
    }
    resourcepath: "/vm/24/utilization"
    resourceid: 24
    utilization: {
      MemConfigured: 25600
      FirstObservedTime: 1435719611
      AllocatedDiskDays: 0.69
      AllocatedMemDays: 0
    }
  }
}
```

```
AvgMemUtil: 0
DiskProvisioned: 0
DiskUtil: 0
AvgCPUTotalUtil: 0
CPUClockSpeed: 2699
AllocatedCPUDays: 0
DiskUsed: 204801
LastObservedTime: 1443603837
NumCPU: 8
}
}
1: {
  properties: {
    SystemVirtType: "VMWARE"
    SystemID: "421a33ea-dcac-3632-02c6-b734c7c014a6"
    SystemName: "CentOS-KVMHost-vm2"
    SystemOSName: "CentOS 4/5/6 (64-bit)"
    SystemHostName: "vpvqa8.ind.sample1.com"
  }
  resourcepath: "/vm/25/utilization"
  resourceid: 25
  utilization: {
    MemConfigured: 20480
    FirstObservedTime: 1435719611
    AllocatedDiskDays: 0.69
    AllocatedMemDays: 0
    AvgMemUtil: 0
    DiskProvisioned: 0
```



```
DiskUtil: 0
AvgCPUTotalUtil: 0
CPUClockSpeed: 2699
AllocatedCPUDays: 0
DiskUsed: 204800
LastObservedTime: 1443603837
NumCPU: 8
}
}
-----
99: {
  properties: {
    SystemVirtType: "KVM"
    SystemID: "ee37967f-0e48-019e-77fc-4175f8578526"
    SystemName: "CentOS-KVM-vm34"
    SystemOSName: ""
    SystemHostName: "vpvqa8.ind.sample2.com"
  }
  resourcepath: "/vm/360/utilization"
  resourceid: 360
  utilization: {
    MemConfigured: 512
    FirstObservedTime: 1435719611
    AllocatedDiskDays: 0.03
    AllocatedMemDays: 0
    AvgMemUtil: 74.41
    DiskProvisioned: 8192
    DiskUtil: 16.43
```

```
AvgCPUTotalUtil: 0
CPUClockSpeed: 3500
AllocatedCPUDays: 0
DiskUsed: 1346
LastObservedTime: 1443603814
NumCPU: 1
}
}
-
}
```

VMs in all Business Groups:

```
{
  "instancecount" : 9,
  "instances" : [ {
    "properties" : {
      "Id" : "Windows",
      "GroupingName" : "VMGROUP",
      "Name" : "Windows"
    },
    "resourcepath" : "/group/VMGROUP/1/utilization",
    "resourceid" : 1,
    "utilization" : {
      "MemConfigured" : 78080.0,
      "CPUCycleTotalUsed" : 1520.33,
      "MemUsed" : 42651.06,
      "CPUMhzCapacity" : 131814.0,
      "MemDays" : 7.74,
      "NumEntities" : 46.0,
```

```
"MemEntlUtil" : 28.25,
"DiskDays" : 231.58,
"AvgCPUTotalUtil" : 1.15,
"CpuDays" : 7.58,
"AvgMemUtil" : 28.25,
"NumCPU" : 51.0
}
},
-----
{
"properties" : {
  "Id" : "BSGDS_data",
  "GroupingName" : "BSGDS",
  "Name" : "BSGDS_data"
},
"resourcepath" : "/group/BSGDS/1/utilization",
"resourceid" : 1,
"utilization" : {
  "Capacity" : 1.178496E7,
  "NumEntities" : 16.0,
  "DiskUsed" : 8781593.0
}
},
-----
{
"properties" : {
  "Id" : "BSGHOST_vmware",
  "GroupingName" : "BSGHOST",
```

```
"Name" : "BSGHOST_vmware"
},
"resourcepath" : "/group/BSGHOST/1/utilization",
"resourceid" : 1,
"utilization" : {
  "MemUsed" : 471616.0,
  "MemConfigured" : 1916457.0,
  "CPUMhzCapacity" : 1455264.0,
  "MemDays" : 325.41,
  "NumEntities" : 13.0,
  "MemEntlUtil" : 2.97,
  "AvgCPUTotalUtil" : 3.81,
  "DiskDays" : 0.0,
  "CpuDays" : 105.0,
  "AvgMemUtil" : 2.97,
  "NumCPU" : 544.0
}
}
```

Note:

CPUDays - It is defined as the product of number of CPUs and number of days the VM is running. For example, if 4 CPUs are used for first 15 days and 8 CPUs are used for next 15 days in a month, the CPUDays is calculated as follows:
 $(4 \times 15) + (8 \times 15) = 180$ CPUDays.

MemDays - It is defined as the product of the amount of memory allocated and number of days the VM is running. For example, if a VM is running with 4 GB of memory for 5 days, the MemDays is calculated as follows: $5 \times 4 = 20$ Memory (GB)-Days.

DiskDays - It is defined as the product of the amount of disk provisioned and number of days the VM is running. For example, if a VM is running with 20 GB of hard disk for 5 days, the DiskDays is calculated as $20 \times 5 = 100$ Disk (GB)-Days.

Using REST APIs to view Snapshots

HPE Cloud Optimizer provides REST APIs to view the details of Snapshots in your virtualized environment.

Resource Model

The REST API to view Snapshots can be invoked at the following levels:

- VM
- Cluster
- Datacenter
- Host

Resource Information

Response formats	JSON
Request Header requirements	Content-Type: application/json Authorization: Basic <Base 64 encoded username:password> (optional, required only when LDAP is enabled)
Supported request methods	GET

Resource URL

The Base URL to get information about Snapshots is:

`https://<IP address or host name>:<port_number>/PV/api/v1`

In this instance,

- `<ip_address or host name>` is the IP address or host name of the HPE Cloud Optimizer server.
- `<port_number>` is the port number of the HPE Cloud Optimizer server. The default port number is 8444.

The Base URL must be appended with appropriate URIs to get information about the Snapshots in the environment. The following table summarizes REST API URL details:

API	URL
To view the details of	<code>https://<ip_address>:<port_number>/PV/api/v1/snapshot</code>

API	URL
all the Snapshots in the virtualized environment	
To view the details of a Snapshot, when you have the name of the Snapshot	<p><code>https://<ip_address>:<port_number>/PV/api/v1/snapshot?name=<snapshot_name></code></p> <p>In this instance, <snapshot_name> is the name of the snapshot.</p>
To view the details of all the snapshots in a VM	<p><code>https://<ip_address>:<port_number>/PV/api/v1/vm/<VM_id>/snapshot</code></p> <p>In this instance, <VM_id> is the ID of the VM.</p>
To view the details of all the snapshots in a Cluster.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/cluster/<cluster_id>/snapshot</code></p> <p>In this instance, <cluster_id> is the ID of the Cluster.</p>
To view the details of all the snapshots in a Datacenter	<p><code>https://<ip_address>:<port_number>/PV/api/v1/datacenter/<datacenter_id>/snapshot</code></p> <p>In this instance, <datacenter_id> is the ID for that datacenter.</p>
To view the details of all the snapshots in a Host.	<p><code>https://<ip_address>:<port_number>/PV/api/v1/host/<host_id>/snapshot</code></p> <p>In this instance, <host_id> is the ID of the host.</p>

JSON Response

The following JSON is returned in the response:

```
{
  "snapshotcount" : 2,
  "snapshotlist" : [ {
    "name" : "2.20_13thJuly2015",
    "description" : "Description of the VM",
    "id" : 7,
    "vmname" : "VM_Name-(xx.xxx.xx.xxx)",
    "creationtime" : 1436795393,
    "VM_OPCCIID" : "VCENTER@!VCENTER_GUEST__12@!localhost@!VCENTER_GUEST",
    "vm_OPCCIID" : "VCENTER@!VCENTER_GUEST__12@!localhost@!VCENTER_GUEST"
  }, {
    "name" : "GartnerDemoState",
    "description" : "Description of the VM",
    "id" : 8,
    "vmname" : "VM_Name-(xx.xxx.xx.xxx)",
    "creationtime" : 1437732934,
    "VM_OPCCIID" : "VCENTER@!VCENTER_GUEST__12@!localhost@!VCENTER_GUEST",
    "vm_OPCCIID" : "VCENTER@!VCENTER_GUEST__12@!localhost@!VCENTER_GUEST"
  } ]
}
```

Appendix

- **CPU Allocation** - Process by which CPU is allocated to run programs and services. The chart shows the amount of CPU allocated for the resources in the domain, in GHz.
- **Memory Allocation** - Process by which virtual memory is allocated to programs and services to access when running. The chart shows the amount of memory allocated to the resources in the domain, in GB.
- **Storage Allocation** - Process by which particular amount of storage space is assigned for specific programs and services. The chart shows the amount of storage allocated for resources in the domain, in GB.
- **CPU Usage** - Indicates the amount of CPU used by the programs and services in the domain, in GHz.
- **Memory Usage** - Indicates the amount of memory used by the programs and services in the domain, in GB.
- **Storage Usage** - Indicates the amount of storage space used by resources in the domain, in GB.
- **CPU Capacity** - Indicates the total CPU, in GHz.
- **Memory Capacity** - Indicates the total memory capacity, in GB.
- **Storage Capacity** - Indicates the total storage capacity, in GB.

Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on PDF version of the Online Help (Cloud Optimizer 3.00)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to docfeedback@hpe.com.

We appreciate your feedback!