

HP Virtualization Performance Viewer

For the Windows® and Linux operating systems

Software Version: 1.20

PDF version of the Online Help

Document Release Date: December 2013

Software Release Date: December 2013



Legal Notices

Warranty

The only warranties for HP 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. HP 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 HP 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 2012-2013 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® is a U.S. registered trademark of Microsoft Corporation.

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 (ey@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: <http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to: <http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login 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: <http://www.hp.com/go/hpsupport>

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:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

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

Contents	4
Chapter 1: Introduction to HP Virtualization Performance Viewer	8
Chapter 2: Administration Overview	9
Data Sources	9
Adding a Data Source	10
Adding Microsoft SCVMM Data Source	11
Adding VMware vCenter Server Data Source	12
Adding KVM Data Sources	13
Adding Xen Data Sources	14
Adding HP aPaaS Data Sources	15
Adding OpenStack Data Sources	17
Removing a Data Source	18
Restarting Collection	18
Data Management	19
Integrations	19
Integrating PM with vPV	19
Integrating BSM with vPV	20
Integrating CSA with vPV	20
License Management	21
User Interface	21
Importing Licenses	21
Time Zones	22
Administrator Tasks	22
Integrating with Microsoft Active Directory and OpenLDAP	24
Configuring LDAP Connection over SSL	26
Obtaining the Server Certificate	27
Importing the Server Certificate	27
Accessing vPV User Interface	27
Configuring Session Timeout Interval	28

Configuring Secure Client Connect Timeout Interval	28
Configuring Java Virtual Machine Memory	28
Daily Maintenance	29
Configuring vPV to use Public Key Infrastructure Authentication	30
Disabling Certificate-based Authentication	32
Chapter 3: Treemap Overview	33
User Interface	33
Using the Treemap	36
Resource Information in Treemap	37
Resources Pane	38
Group by and Color by	38
Using the Drill Down	41
Searching Resources	41
Filtering Data	42
Micro Charts	42
Viewing Micro Charts	43
Use case: Viewing Data in the Treemap	43
Chapter 4: Optimization and Placement Overview	45
Personas	45
Optimization Overview User Interface	45
Terminologies	49
Summary for Datacenter	50
Summary for Cluster	51
Summary for Host	51
Summary for VM	54
Summary for Datastore	56
CPU Sizing	58
Memory Sizing	59
CPU and Memory Reservation	60
Using Placement	61
Example: Using Capacity Usage Trends and Optimization Option to Manage the Resource Utilization in your Environment	62

Chapter 5: Forecast	66
Using Forecast to plan for resources	68
Forecasting CPU Utilization	69
Forecasting Memory Utilization	69
Forecasting Disk Utilization	70
Chapter 6: Monitoring the Infrastructure as a Cloud Service Automation (CSA) user	72
Chapter 7: Real Time Guest OS Drilldown for Advance Troubleshooting	73
Launch the Tool	73
User Interface Details of Guest OS DrillDown	73
Graphs	74
Tabs	75
Options	80
Using Guest OS Drilldown to Troubleshoot the Resource Utilization in your Environment	81
Chapter 8: Performance Graphing Overview	82
Overview of Graphs	82
Metrics	82
Icon Glossary	93
Workbench Overview	93
Configuration Items	94
Favorites	94
Performance Pane	94
Options for Drawn Graphs	95
Table Graph Window	97
Using Table Highlights	99
Using Table Filters	100
Graph Export Dialog Box	101
Date Range Panel	101
Drawing Graphs	103
Saving as Favorite	104
Deleting a Favorite	104
Drawn Graphs Functions	105

Reports Overview	106
Viewing Reports	106
Types of Reports	107
Chapter 9: Troubleshooting vPV	109
Chapter 10: Frequently Asked Questions	121
We appreciate your feedback!	129

Chapter 1: Introduction to HP Virtualization Performance Viewer

HP Virtualization Performance Viewer (vPV) is a web-based tool that helps you monitor the resources in virtualized and cloud environment. After installing vPV in your environment, you can add data sources and start using the tool to monitor the resources. You can use the performance data to troubleshoot problems and to plan and utilize the resources in your environment better.

Using vPV, you can monitor resources in the following data sources:

- VMware vCenter Servers
- Microsoft System Center Virtual Machine Manager (SCVMM)
- KVM
- Xen
- OpenStack
- HP aPaaS


The key features of vPV are as follows:

- **Treemap** - Treemap provides you a graphical representation of the resource utilization in your monitored environment. For more information on using the Treemap, see ["Using the Treemap " on page 36](#).
- **Optimization and Placement** - Helps you to utilize and allocate your resources properly. For more information, see [" Chapter 4: Optimization and Placement Overview" on page 45](#)
- **Forecast**: View forecast data of your resource utilization for next 90 days. For more information, see ["Chapter 5: Forecast"](#)
- **Performance Graphing** - Performance Graphing helps you to draw predefined graphs or graphs for metrics based on your resources, in the Workbench page. For more information on using the graphs, see ["Chapter 8: Performance Graphing Overview " on page 82](#).
- **Reporting** - Reports show you performance and status data corresponding to your resources, in the Reports tab. For more information on the available reports, see ["Reports Overview" on page 106](#)
- **Administration** - Using the Administration page, you can manage vPV and perform Administrator tasks, such as, adding and removing data sources. For more information, see ["Chapter 2: Administration Overview" on page 9](#).

Note: vPV supports the VMware vCenter Server versions 4.1, 5.0, 5.1, and 5.5 and Microsoft SCVMM version 2012.

Chapter 2: Administration Overview

After installing (vPV), you can add the data sources in your network, start monitoring resources, and view the resource utilization.

When you log on to vPV for the first time, by default, the Administration page opens. Add a data source and start using the product. To launch the Administration page from the vPV home page, click  in the Options pane.

Note: Only a user with Administrator privileges can add or remove data sources in the Administration page. The system checks for the user privileges while you log on and if you do not have the sufficient privileges, the Administration page is available only in Read-only mode. In the Read-only mode you cannot add or remove data sources.

The following table lists the tabs available in the Administration page and the corresponding functions.

Name	Description
Data Sources	<p>Use this tab to add or remove data sources. By default, this tab is selected when you launch vPV and launch the Administration page for the first time. For more information, see "Data Sources" below.</p> <p>You can share your queries or feedback about the product at the HP Software Community. To access the Community portal, click the HP Software Community link on the Data Sources tab of the Administration page.</p>
Data Management	<p>This tab displays the data collection and retention information for vPV. For more information, see "Data Management" on page 19.</p>
Integrations	<p>Use this tab to integrate and use vPV with other HP products - HP Business Service Management and HP Performance Manager. For more information, see "Integrations".</p>
License Management	<p>This tab displays information about your vPV licenses. You can also activate your evaluation license and import permanent licenses using this tab. For more information, see "License Management" on page 21.</p>

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

Data Sources

Using the **Data Sources** tab you can select a domain, add a data source to be monitored, remove any added data source, and also see the list of data sources that are monitored by vPV. After adding a data source, the status and details of the data source appear in the Connections List at the bottom of the page.

The following table lists the elements available in the Connections List.

Name	Description
IP/Host Name	Lists the IP addresses or host names of all the data sources added to vPV for monitoring.
Domain	Domain of the data source added for monitoring.
User Name	User name corresponding to the added data source.
Instance Count	Number of instances associated with each monitored data source. 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 status messages, see "Data Collection Status Messages" on page 117
Last Collection Time	Shows the date and client time zone against each data source, when data was last collected. Note: This field does not show any value till the first data collection is complete.

Note: vPV data source collection works only when sum of instances from all data sources 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 data sources 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.

Following are the tasks that can be performed from the Data Sources tab:

- ["Adding a Data Source" below](#)
- ["Removing a Data Source" on page 18](#)

Adding a Data Source

To start monitoring the utilization of resources in your virtualized environment, you must add data sources to vPV. When you add the data source to be monitored, vPV 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.

Note: At times, when you add a data source, one more row gets created in the Connections list. This does not impact the data collection.

Using vPV you can add and monitor the following data sources:

- [Microsoft SCVMM](#)
- [VMware vCenter](#)
- [KVM](#)
- [Xen](#)
- [HP aPaaS](#)
- [OpenStack](#)

Adding Microsoft SCVMM Data Source

Prerequisites

- .NET framework 3.0 or above
- Microsoft SCVMM 2012
- SCVMM Admin console must be installed

Note: You must have Administrator rights in SCVMM to run the scripts.

If the password for the Administrator user is changed, remove the collector and add it again to ensure continuous collection. For more information on removing data sources, see ["Removing a Data Source" on page 18](#)

To add a Microsoft SCVMM data source, follow these steps:

1. Click **Data Sources** tab in the Administration page.

If you are opening the Administration page for the first time after starting vPV, the Data Sources tab is selected by default.

2. Select **Microsoft SCVMM** from the **Domain** drop-down list.
3. Click **Download Collector** to download the `PV_SCVMMCollectorScript.zip` file.
4. Extract the file on the Microsoft SCVMM host.
5. Run the `Start-Collector.bat` file.

Note: You must manually copy the `PV_SCVMMCollectorScript.zip` file to the target Microsoft SCVMM host. If you download the file directly from the browser to the Microsoft SCVMM host, there might be some permission issues.

Ensure that you use the FQDN for your browser when downloading the file.

Adding VMware vCenter Server Data Source

vPV supports the VMware vCenter Server versions 4.1, 5.0, 5.1, and 5.5.

To add a VMware vCenter Server data source, follow these steps:

1. Click **Data Sources** tab in the Administration page.

If you are opening the Administration page for the first time after starting vPV, the Data Sources tab is selected by default.

2. Select **VMware vCenter** from the **Domain** drop-down list.
3. Type the IP address or host name of the data source in the **vCenter IP/Host Name** field.

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

4. Type the user name corresponding to the specified data source in the **User Name** field.
5. Type the password corresponding to the specified data source in the **Password** 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 Connections list displays the added data sources.

Note: To add a VMware vCenter Server as a data source in vPV, the user must have the following roles in addition to the Read-only role:

- Set the **Browse Datastore** role available under **Datastore**
- Set the **Validate session** role available under **Sessions**

Also, you must **Enable Statistics** on the VMware vCenter Server. For more information, see the *VMware documentation*.

¹Fully Qualified Domain Name

Adding KVM Data Sources

Prerequisites

- Install **Libvirt** and all its dependencies on the vPV 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 flavors on which you can install Libvirt.

Linux Types	OS Versions
CentOS	6.2 and 6.3
RHEL	6.2 and 6.3
Ubuntu	12.04

- Install **Expect** and all its dependencies on the vPV server.

Expect is a tool that communicates with interactive programs, such as, ssh. If Expect is not installed on the system, make the following changes:

- Open the `parm` file from the location `/var/opt/perf`.
- In the `parm` file, set the `ls_collection` parameter to `native`.
- Save the file and exit.
- Open the `automate_ssh.exp` script from `\opt\OV\bin\`.
- Run the following commands from the script:

```
ssh-keygen -t rsa
```

```
ssh-copy-id
```

To add a KVM data source, follow these steps:

1. Click **Data Sources** tab in the Administration page.

If you are opening the Administration page for the first time after starting vPV, the Data Sources tab is selected by default.

2. Select **KVM** from the **Domain** drop-down list.
3. Type the IP address or host name of the data source in the **IP/Host Name** field.
4. Type the user name corresponding to the specified data source in the **User Name** field.

5. 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.

6. Click **Add**.

The Connections list displays the added data source.

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

Adding Xen Data Sources

Prerequisites

- Install Libvirt and all its dependencies on the vPV server.

Libvirt is an open source management tool to manage virtualized platforms such as Linux, KVM, Xen, and so on. Libvirt can be installed on following versions of the different Linux flavors:

Linux Types	OS Versions
Ubuntu	12.04
SLES	11 Service Pack 2

- Install Expect and all its dependencies on the vPV server.

Expect is a tool that communicates with interactive programs such as ssh. If Expect is not installed on the system, make the following changes:

- a. Open the `parm` file from the `/var/opt/perf` folder.
- b. In the `parm` file, set the `ls_collection` parameter to **libvirt**.
- c. Save the file and exit.
- d. Open the `automate_ssh.exp` script from `\opt\OV\bin\`.
- e. Run the following commands from the script:

```
ssh-keygen -t rsa
```

```
ssh-copy-id
```

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 data source, follow these steps:

1. Click **Data Sources** tab in the Administration page.

If you are opening the Administration page for the first time after starting vPV, the Data Sources tab is selected by default.

2. Select **Xen** from the Domain drop-down list in the **Data Source** tab.
3. Type the IP address or host name of the data source in the **IP/Host Name** field.
4. Type the user name corresponding to the specified data source in the **User Name** field.
5. 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.

6. Click **Add**.

The Connections list displays the added data source.

Adding HP aPaaS Data Sources

vPV supports the HP aPaaS version 2.10.

Note: The username or password of HP aPaaS must not contain any of the following special characters -/, \, %, and ?.

Prerequisites

Before adding HP aPaaS data sources, you must configure the `collectd` component of HP aPaaS.

Note: vPV does not show any collection errors if `collectd` is not configured or it is not running. But there will be no HP aPaaS node data collected.

To configure the `collectd` component of HP aPaaS:

1. Log on to the HP aPaaS host as root user.
2. Open `collectd.conf` from `/etc/collectd`.
3. In `collectd.conf`, search for the `<Plugin write_http>` tag.

4. Remove **#** before the `<Plugin write_http>` tag.
5. Add the following contents to the `<Plugin write_http>` tag:

```
<Plugin write_http>

<URL "http://<IPaddress/server_name>:<port_
number>/PV/collectDListener?target=<aPaaS_target_url>">

Format "JSON"

</URL>

</Plugin>
```

where,

- `<IPaddress/server_name>` is the IP address and server name of the vPV server respectively.
- `<port_number>` is the port which vPV uses. The default port number is 8081.
- `<aPaaS_target_url>` is the URL of the HP aPaaS target.

For example

```
<Plugin write_http>

<URL "http://12.220.15.90:8081/PV/collectDListener?target=example.apaas-
abc1.local">

Format "JSON"

</URL>

</Plugin>
```

6. Restart collectd service using the following command:

```
service collectd restart
```

To add a HP aPaaS data source, follow these steps:

1. Click **Data Sources** tab in the Administration page.

If you are opening the Administration page for the first time after starting vPV, the Data Sources tab is selected by default.

2. Select **HP aPaaS** from the Domain drop-down list in the **Data Source** tab.

3. In the **HP aPaaS URL** field, type the URL of the machine where HP aPaaS is installed. For example: `https://example.apaas-abc1.local`.
4. Type the email corresponding to the specified data source in the **Email** field.
5. Type the password corresponding to the specified user name in the **Password** 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 Connections list displays the added data source.

Adding OpenStack Data Sources

Prerequisite

Healthnmon must be installed on the cloud controller before you start adding OpenStack data sources. Healthnmon collects utilization metrics which are then collected by vPV. To install Healthnmon, go to the following Web site: <https://github.com/stackforge/healthnmon>

To add an OpenStack data source, follow these steps:

1. Click **Data Sources** tab in the Administration page.

If you are opening the Administration page for the first time after starting vPV, the Data Sources tab is selected by default.
2. Select **OpenStack** from the Domain drop-down list in the **Data Source** tab.
3. Type the IP address or host name of the target data source in the **Target IP/Host Name** field.
4. The **Port** field displays the default port number for OpenStack.
5. Type the user name corresponding to the specified data source in the **User Name** field.
6. Type the password corresponding to the specified user name in the **Password** field.
7. Type the authentication token that you receive upon installation of OpenStack on the host in the **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 Connections list displays the added data source.

Note: If a data source is down, vPV does not collect data for that data source.

Removing a Data Source

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

To remove a data source from the list of monitored connections, follow these steps:

Note: For a Microsoft SCVMM data source, run the Stop-Collector.bat file first.

1. Select the data source you want to remove from the list of added connections.

You can select multiple connections by pressing the **Ctrl** key and clicking the required connections. The **Remove** button is enabled.

2. Click **Remove**.

The system displays a dialog box asking for confirmation to remove the data source.

3. Click **OK** to remove the data source.

The system displays a confirmation message and no longer monitors the removed data source. The data source is also removed from the Connections list.

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

Restarting Collection

To restart the data source collection:

1. Click **Data Sources** tab in the Administration page.

If you are opening the Administration page for the first time after accessing vPV, the Data Sources tab is selected by default.

2. Select the data source that you want to restart.
3. Click **Restart Collection**.

Data Management

The Data Management tab provides you information about the data retention by vPV. The tab 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*.

Apart from the data retention information, the tab also displays the following information:

- Current database size
- Available Free space

Integrations

The Integrations page helps you to integrate vPV and use it with products like HP Performance Manager (PM) and HP Business Service Management (BSM). When you integrate PM, you can launch the PM graphs against VMs from the Treemap.

Note: If you are not able to view all elements on the Integrations tab, maximize the tab.

The Integrations feature is available only with the Evaluation and Permanent licensed versions of vPV.

Integrating PM with vPV

You can draw URL based graphs for VMs with HP Operations agent installed on them for the vCenter and SCVMM data sources.

For the vCenter data sources, the VMware tools must be installed and running on the VMs and the IP of the VM must be pingable from the PM server. The integration does not work properly if authentication is enabled on PM. If authentication is enabled, you must pass the username and password in the URL during the integration. You can launch only one graph at a time using the PM integration.

To integrate a PM system and draw graphs against VMs, follow these steps:

1. Type a name for the integration in the **Name** field.
2. Type the URL for the PM integration in the **URL** field.
3. Click **Add/Update**. The integration is added to the Integrations list.

Note: To remove an added PM integration URL, select the URL and click **Remove**. The system displays a confirmation message, click **Yes**. To remove multiple URLs, press the CTRL key and select the URLs.

Integrating BSM with vPV

To integrate BSM with vPV, follow these steps:

1. 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.
2. Type the key in the BSM Token Creation Key field.

Note: Use the token only if authentication is enabled in vPV

3. Click **Save**.
4. Follow the instructions that appear in the dialog box.

Integrating CSA with vPV

To integrate CSA with vPV, follow these steps:

1. Type the URL for the integration in the **CSA URL** field.
2. Type the user name in the **User Name** field.
3. Type the password in the **Password** field.
4. Click **Save** and log out from vPV.

Note: After the configuration, a new CSA user is created. The CSA user can check the performance of the VMs by logging on as CSA user.

For more information, see "[Chapter 6: Monitoring the Infrastructure as a Cloud Service Automation \(CSA\) user](#)".

License Management

The License Management tab provides information about your vPV licenses. You can also use this tab to start using your Evaluation License as well as import a vPV license. For more information on importing licenses, see ["Importing Licenses" below](#).

For more information on the difference between the licenses and the features available, see *HP Virtualization Performance Viewer Installation Guide*.

User Interface

The following table lists the sections available in the License Management tab.

Section	Description
License Status	Displays information regarding the installed licenses. The two tables available are: <ul style="list-style-type: none"> • Active License Details • Installed Licenses
Manage License	To import permanent licenses.

The **Active License Details** table lists the following information about your current vPV License:

- Type of license installed
- Date of license expiry and number of days remaining until license expiry
- Number of instances that can be monitored, based on the installed license
- Number of instances monitored currently

The **Installed Licenses** table lists the following information regarding the active and all installed licenses for vPV:

- All active and installed licenses.
- Capacity or number of instances available corresponding to each license.

Note: The Installed Licenses table is not available when your current active license is Community license.

Importing Licenses

To import licenses for vPV, follow these steps:

1. Go to the **License Management** tab in Administration page.
2. Type your license key in the **License Key** field of the Manage License section.
3. Click **Import License**. The License Status section refreshes to display the details of the license imported.

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

Time Zones

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

If you are using vPV to monitor Microsoft SCVMM data sources and if the vPV server and Microsoft SCVMM servers are in the same time zone, make sure that the time is synched.

Administrator Tasks

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

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

- ["Integrating with Microsoft Active Directory and OpenLDAP"](#)
- ["Configuring Session Timeout Interval" on page 28](#)
- ["Configuring Secure Client Connect Timeout Interval" on page 28](#)
- ["Configuring Java Virtual Machine Memory" on page 28](#)
- ["Configuring vPV to use Public Key Infrastructure Authentication"](#)

Conventions Used

The following conventions are used in the sections when referring to the location of files on the vPV server system.

- *<install_dir>* - The directory where vPV is installed. The default location is */opt/OV*. 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 HP Software products are stored. The default location is */var/opt/OV*. For Archive Extractor version of the product, the location is */data*.
- *<bin_dir>* - The directory where the binaries reside. The default location is */opt/OV/bin*. For

Archive Extractor version of the product, the location is same as the `<install_dir>` location.

- `<systemname>` - The name of the system on which the vPV server is running.

When accessing vPV through a web browser, you must replace this variable with the actual system name or IP address of the system.

OVPConfig.ini Parameters

The following table lists the parameters, which must be specified in the OVPConfig.ini file to configure vPV.

Parameter	Description
TRACELEVEL	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. 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.
GRAPH_ AUTOREFRESH_ RATE	You can use this parameter to specify the interval at which vPV 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 29 .
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 vPV 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 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. Note: The value selected from the user interface menu option overrides the setting specified for this parameter.
SECURE_ CLIENT_ CONNECT_TIM EOUT	You can use this parameter to specify the timeout period for which secure communication HTTPS protocol is used for connecting to the vPV server. The default timeout period is one second. For more information, see "Configuring Secure Client Connect Timeout Interval" on page 28 .

Parameter	Description
RTV_GRAPH_DATAPOINTS	You can use this parameter to configure the maximum number of data points in the graphs drawn using Workbench.

Integrating with Microsoft Active Directory and OpenLDAP

vPV supports authentication using Microsoft Active Directory and OpenLDAP. To configure vPV to use the Microsoft Active Directory server or OpenLDAP, perform the following steps:

1. Open the `OVPMconfig.ini` file available in the `<data_dir>/conf/perf` directory.
2. Add the information in the `[LDAP]` section of the `OVPMconfig.ini` file by removing the comment marks “;”, from the example lines and replacing information with the values that match your installation.

Example for Microsoft Active Directory

```

PVGROUP=<User group for PV> which is Active Directory user group.

PVADMIN=<A user account present in Active Directory and will be used as PV
Administrator, example PV_Admin. This user will have the Administrator privileges for
vPV and must be a part of the PVGROUP>.

LDAPHOST=<IP/FQDN of Active Directory Server, example Sample.abc.com>

SEARCHBASE=<The server root under which vPV will search for users>. For example, if
the domain name is abc.def.com then SEARCHBASE will be
DC=<abc>,DC=<def>,DC=<com>

USERSEARCHQUERY=(&(objectclass=user)(SAMAccountName=$USERID$))

DOMAIN=<abc.def.com>

SSL_KEYSTORE=<Absolute path of keystore file>, if enabled.

USE_SSL=true/false. This is to switch the SSL ON or OFF for the authentication of the
users. This is only in the vPV context. For CSA users, using SSL for LDAP authentication
is based on the CSA configuration for LDAP.

```

Example for OpenLDAP simple configuration

```

In this scenario, all vPV users are under the same organization (OU).

PVADMIN=<A user account used as PV Administrator, example PV_Manager. This user

```


will have the Administrator privileges for vPV>.

LDAPHOST=<IP/FQDN of OpenLDAP Server, example Sample.abc.com>

SEARCHBASE=<The server root under which vPV will search for users>. For example, if the domain name is my-domain.com then SEARCHBASE=OU=MyOrganization,DC=my-domain,DC=com. MyOrganization is an organization created in OpenLDAP.

*USERSEARCHQUERY=CN=\$USERID\$. This field is to specify the attribute to identify the user. For example if "CN" is the attribute, then
USERSEARCHQUERY=CN=\$USEERID\$. If this is UID then
USERSEARCHQUERY=UID=\$USERID\$*

DOMAIN=<my-domain.com>

LDAPTYPE=OpenLDAP. The default value is AD.

SSL_KEYSTORE=<Absolute path of keystore file>, if enabled.

USE_SSL=true/false. This is to switch the SSL ON or OFF for the authentication of the users

Example for OpenLDAP Group configuration

In this scenario, vPV users are part of different Organizations in LDAP.

PVGROUP=<User group for PV> which is OpenLDAP user group.

PVADMIN=<A user account used as PV Administrator, example PV_Manager. This user must be a part of PVGROUP and will have the Administrator privileges for vPV>.

LDAPHOST=<IP/FQDN of OpenLDAP Server, example Sample.abc.com>

SEARCHBASE=<The server root under which vPV will search for users>. For example, if the domain name is my-domain.com then SEARCHBASE=DC=my-domain,DC=com

*USERSEARCHQUERY=CN=\$USERID\$. This field is to specify the attribute to identify the user. For example if "CN" is the attribute, then
USERSEARCHQUERY=CN=\$USEERID\$. If this is UID then
USERSEARCHQUERY=UID=\$USERID\$*

DOMAIN=<my-domain.com>

LDAPTYPE=OpenLDAP. The default value is AD.

BIND_DN=CN=user1, OU=test,DC=my-domain,DC=com. This user must have read

access to the SERACHBASE,PVGROUP and attributes.

BIND_DN_PASSWORD=Encrypted password of above DN (Note: Use XPL to obtain the encrypted password).

SSL_KEYSTORE=<Absolute path of keystore file>, if enabled.

USE_SSL=true/false. This is to switch the SSL ON or OFF for the authentication of the users

Note: To generate the encrypted password:

1. Use the command `#/opt/OV/bin/pvconfig`

Syntax : `/opt/OV/bin/pvconfig -en <Any password>`

Example: `#/opt/OV/bin/pvconfig -en password`

An example of encrypted password - `dAGZEfcZEPiQxXNilr85Cxc81jsomV8v`

2. Copy the password as `BIND_DN_PASSWORD` in

`/var/opt/OV/conf/perf/OVPMconfig.ini` under `[LDAP]` name space.

3. Restart ovtomcatB using the command `# /opt/OV/bin/ovc -restart ovtomcatB.`

Configuring LDAP Connection over SSL

You can configure LDAP connection with SSL if you want a secure connection between the vPV server and the LDAP server.

To configure LDAP connection over SSL (LDAPS), perform the following steps:

1. Open the `OVPMconfig.ini` file available in the `<data_dir>/conf/perf` directory.
2. Add the information in the `[LDAP]` section of the `OVPMconfig.ini` file by removing the comment marks “;”, and replacing information with the values that match your installation.

SSL_KEYSTORE=<data_dir>\conf\perf\jssecacerts. This is the location of keystore that stores the LDAP server certificate.

Set `USE_SSL=true`

This enables the communication between LDAP and vPV through LDAPS. Otherwise, the communication takes place through LDAP.

3. Restart vPV using the following command:

`pv restart`

Obtaining the Server Certificate

You must add the Microsoft Active Directory server SSL certificate to the list of accepted certificates used by the vPV server. To add the certificate, export the certificate by running the following command on the Microsoft Active Directory server:

`certutil -ca.cert <sample.crt>`, where `sample.crt` is the name of the SSL certificate that you want to export to the vPV server.

Importing the Server Certificate

You must import the Microsoft Active Directory server certificate into the keystore for SSL enabled communication between the vPV and Microsoft Active Directory.

The keystore file, for example, `jssecacerts` is located in the `<data_dir>/conf/perf` directory.

Note: The `keytool` command creates the keystore file if it is not present in the directory.

To import the server certificate, perform the following steps:

1. Go to `<data_dir>/conf/perf` directory.
2. Copy the exported Microsoft Active Directory server SSL certificate and paste the certificate in the directory.
3. Run the following command:

```
/opt/OV/nonOV/jre/b/bin/keytool -importcert -keystore jssecacerts -file /root/cacert.pem. The
keytool prompts you for a password. The default password is changeit.
```

4. Select yes to confirm the key import when prompted with `Trust this Certificate?[no]:`
yes

Accessing vPV User Interface

To launch the Microsoft Active Directory enabled vPV, you can use the following URL:

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

If you launch the Microsoft Active Directory enabled vPV 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 HP Virtualization Performance Viewer 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 vPV server over SSL. However, if you initiate vPV using the URL: **https://<system>:<port>/PV**, vPV continues to run in the secured mode.

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 `OVPMconfig.ini` file available in the `<data_dir>/conf/perf` directory.
2. Specify an interval for the `SESSION_TIMEOUT` parameter in minutes. For example, set it to 30 minutes.
3. Save the file.
4. Restart vPV.

Configuring Secure Client Connect Timeout Interval

As a vPV administrator, you can set the expiration time or timeout interval at which vPV can connect to the server through HTTPS communication protocol. When vPV 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 `OVPMconfig.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 for 10 seconds.

3. Save the file.
4. Restart vPV.

Note: When vPV 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 vPV to stop accepting requests by using the `JVM_MIN_MEMORY` parameter.

By setting this parameter to a specific value, vPV 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 vPV server has run out of memory.

If the request is for XML graphs, vPV 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 `OVPMconfig.ini` file.

To specify `JVM_MIN_MEMORY`, perform the following steps:

1. Open the `OVPMconfig.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, vPV 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 vPV.

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.

Daily Maintenance

vPV performs maintenance task on a daily basis for the following:

- To reclaim disk space for files and directories that are no longer needed
- To close any data sources 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 `OVPMconfig.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 vPV.

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 for 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.

Configuring vPV 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. vPV 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 vPV to use the client certificates, users can log on to vPV with the X.509 certificates without having to manually type in user names and passwords.

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

The prerequisites for using vPV with CAC are as follows:

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

To configure vPV 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 vPV 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 vPV.

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, vPV 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 vPV 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
```

vPV 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
```


Chapter 3: Treemap Overview

In the vPV home page, you can see the utilization of the resources for the added data sources. For more information on adding a data source, see ["Adding a Data Source" on page 10](#). You can perform the following tasks on this page:

- 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 5 summarization intervals
- List up to five most and least utilized resources
- Highlight up to five most and least utilized resources

User Interface



The following table lists the features on the vPV home 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" on page 38 .
Options Pane	Contains the fields and filter based on which you can view data in the Treemap. It also contains the Micro charts feature.
Treemap	Displays the resource utilization of the resource type you selected in the Resources pane. For more information, see "Using the Treemap" on page 36 .











Note: The minimum screen resolution required for viewing the vPV interface is 1280x768.






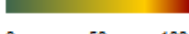

Options Pane

The Options Pane provides you the fields and values based on which you can view data in the Treemap.

- 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.


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" on page 41 .
Remove Highlights ()	Clears the text box corresponding to the Search feature and also clears the Treemap highlighting resulting from the previous search.
Workbench ()	Click to open the Workbench page. For more information, see "Chapter 8: Performance Graphing Overview " on page 82
Capacity Usage Trends 	Click to open the Optimization Summary page. For more information, see "Chapter 4: Optimization and Placement Overview" on page 45 .
Find VM Placement 	Click to open the VM Placement page. For more information, see "Using Placement" on page 61 .
Forecast 	Click to open the Forecast page. For more information, see "Chapter 5: Forecast" on page 66 .
Export as PDF ()	Click to export all the data in the page, including the Treemap, as a PDF document.
Administration ()	Click to open the Administration page. For more information, see "Chapter 2: Administration Overview" on page 9
Help ()	Click to open the context-sensitive help for the page.
Troubleshooting Tips ()	Click to open the help page listing the tips to troubleshoot vPV.
About HP Virtualization Performance Viewer ()	Click to view information about HP Virtualization Performance Viewer.
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" on page 38 .
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.

Field	Description
Summarized Every	Displays the time interval configured for collecting data. The time interval varies depending on the data source 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.
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" on page 42.
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> <div data-bbox="516 1207 966 1633" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p>Data as of 4/26/13 11:25</p> <p>Summarized every 5 minutes</p> <hr/> <p>Size by Available GBs </p> <p>Color by Usage  </p> <p>Filter </p> <p style="text-align: center;"></p> <p style="text-align: center;">0 50 100</p> <p>Spectrum Transition Value 75 </p> <hr/> <p>Hide labels <input type="checkbox"/></p> </div> <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.

Field	Description
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 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.

Note: If there is any problem in the data collection, then the following message will appear at the bottom of the Options pane:



Click the message to go to the **Data Sources** tab on the Administration page and resolve the error.

Using the Treemap

The Treemap provides a visual representation of resources of the data sources you are monitoring. These resources differ based on the type of data source being monitored. You can use the various options in the **Options** pane to filter your view in the Treemap. From the Treemap, you can view the data specific to a resource by right-clicking the resource. 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.
- **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 data source.
- 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.

Note: In HP aPaaS, if an application belongs to a group, but is not a part of any user, then the application is not displayed in vPV.

Resource Information in Treemap

When you move the pointer over resources in the Treemap, you can view information corresponding to the specific resources in a pop-up window. 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.

Note: If a single HP aPaaS user belongs to multiple groups, vPV displays the user information in only one of the groups.

For example, if user U1 belongs to groups G1 and G2, then vPV displays U1 as part of either G1 or G2 but not both the groups.

Related Topics

- ["Using the Drill Down" on page 41](#)

- ["Searching Resources" on page 41](#)
- ["Use case: Viewing Data in the Treemap" on page 43](#)

Resources Pane

The Resources Pane on the top of the vPV 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 also. 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 data sources.

Domain	Resource	Group By	Color By
vSphere	Datacenters	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
	Data Stores	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	

Domain	Resource	Group By	Color By
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
	Virtual Machines	Tenants	<ul style="list-style-type: none"> • CPU Usage • Memory Usage • Disk Usage
	Hosts	Cloud	<ul style="list-style-type: none"> • CPU Usage • Memory Usage
HP aPaaS	Target	None	<ul style="list-style-type: none"> • CPU Usage • Memory Usage
	Nodes	None	<ul style="list-style-type: none"> • CPU Percent • Memory Percent
	User Groups	None	Memory Usage
	Users	User Groups	Memory Usage
	Applications	<ul style="list-style-type: none"> • Group • User 	<ul style="list-style-type: none"> • CPU Usage • Memory Usage • Disk Usage

Note:

In the vPV 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 Drill Down

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.

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.


Example: 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 data stores, the resources corresponding to the search appear highlighted in the Treemap view for data stores.

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 vPV help the user to quickly view trends of the resources being monitored. These micro charts are graphs and correspond to 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 vPV interface lists the top or bottom resources. When you select **Top** or **Bottom**, vPV 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 38](#)

Use case: Viewing Data in the Treemap

This section shows you 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 data stores that are utilizing 80% to 90% of their storage space and view the storage virtual machines associated with these three data stores.

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

1. Log on to the HP Virtualization Performance Viewer home page.
2. In the Resources Pane, click **Data Stores**. The treemap refreshes to show all the data stores corresponding to the VMware vCenter Servers added.


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 data stores that have the storage usage between 80% and 90%.

4. Select the **Top** option. The names and graphs corresponding to the top five data stores that have storage usage in 80% to 90% range are listed.
5. Select **Highlight**. The Treemap highlights the top five utilized data stores in the selected range.


6. Move the pointer over the highlighted data stores to view the exact percentage utilization of the top three data stores.
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 4: Optimization and Placement Overview

vPV provides you the following features to help tune your virtualized infrastructure to maximize efficiency and resource utilization.

Optimization

The optimization feature provides a cohesive view of the resource capacity, usage, allocation, remaining capacity, and opportunity to use the resources more effectively. The reports are based on the utilization and consumption of the resources and focus primarily on reclaiming and adding resources. By using these reports, you can tune your infrastructure for maximum efficiency and

reclaiming wasted resources. To access the feature, click  from the vPV home page. Alternatively, you can also right click an entity from the Treemap, and select **Capacity Usage Trends**.

Placement

The placement feature provides you recommendations about where the new VMs can be added in the environment. The proper allocation of VMs helps in balancing resource utilization across datacenters.

Key benefits of optimization and placement features are as follows:

- Sizing Recommendation - Based on the historical resource usage and demand trends and available capacity, vPV recommends ways to optimize your current resources without impacting the service levels or performance of the product.
- Forecast - You can forecast the resource requirements based on the historical resource usage, so that you can plan for future hardware requirements for your environment.
- General overview of the Virtualization infrastructure - You can identify over-utilized and under-utilized resources, idle VMs, and resolve incorrect VM sizing problems.


For more information on using the feature, see "[Using Placement](#)" on page 61.

Personas

Typical users of this feature are:

- Capacity Planner
- IT Administrator

Optimization Overview User Interface

When you click  from the vPV home page, the **Overview** page for the Optimize feature opens. The following section provides information on the user interface elements on the page.

Toolbar

Drop down items are available on the top of the page. The items available are as follows:

- Home - Click to go back to the Overview page.
- Analysis - Drop down shows the following reports:
 - **VMs Sorted by Reservation** - Lists the VMs having CPU or Memory Reservation. Click to view the ["CPU and Memory Reservation" on page 60](#) page.
 - **VM CPU Sizing** - Lists the VMs with its recommended CPU size, based on historical resource usage. Click to view ["CPU Sizing" on page 58](#).
 - **VM Memory Sizing** - Lists the VMs with its recommended Memory size, based on historical resource usage. Click to view ["Memory Sizing" on page 59](#).
- Tools - Click to open the Workbench page in the current context. You can see all the metrics, quick graphs, and reports for the CI in context. For more information, see ["Chapter 8: Performance Graphing Overview " on page 82](#)

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 ["Terminologies" on page 49](#).

The charts available are as follows:

- **CPU** - Shows the overall CPU allocation, usage, and capacity of the resources.
- **Memory** - Shows the overall Memory allocation, usage, and capacity of the resources.
- **Storage** - 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 on 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 Datacenters in the domain.
- Number of clusters in the domain.
- Number of hosts in the domain.
- Number of powered on VMs in the domain.
- Number of Datastores in the domain.

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:

- Datacenters
- Clusters
- Hosts
- DataStores

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 on the resource attribute name to view the data in ascending or descending order.

Filtering 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	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. When you click a resource name, the page refreshes to show the summary page for the selected resource.
Remaining Capacity	Number of VMs that can be added or created in the Datacenter or Cluster, in addition to the existing ones. The data stores 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.

Attribute	Description
CPU Available for Allocation (GHz)	Shows the total CPU available, in GHz, corresponding to the resource. The value is calculated based on the following formula: 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)	Shows the amount of CPU that can be reclaimed, in GHz, corresponding to the resource. 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.
Memory Available for Allocation (GB)	Shows the total Memory available, in GB, corresponding to the resource. The value is calculated based on the following formula: 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 (GB)	Shows the amount of memory that can be reclaimed, in GB, corresponding to the resource. When you click the attribute value, the page refreshes to display the Memory Sizing page. For more information, see Memory Sizing.

Terminologies

- **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

¹By default, this value is considered as 20%. This means that vPV considers 80% of CPU Util% as healthy.

² $\text{CPUCyclesTotalUsed/TotalCPU}$

³By default, this value is considered as 20%. This means that vPV considers 80% of Memory Util% as healthy.

⁴ $\text{MemPhysUtil/TotalMem}$

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.

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 46.

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 48.

The Resource Information table displays data corresponding to clusters, hosts, and datastores for the particular datacenter. You can click on the resource names and drill down to view the next level of detailed 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 46](#).

The tabular data on the right side of the page displays the domain name, datacenter name, and a count of the hosts, VMs, and data stores 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 the previous page](#).

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 48](#).

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 Domain	The virtualization domain to which the host belongs.
Cluster Name	Name of the cluster to which the host belongs. Click on the cluster name to view the "Summary for Cluster" on page 51 .
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 50 .
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	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 Reserved (GB)	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.

Item Name	Description
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, and 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 Configured
- Memory Reserved (GB)
- CPU Shares Priority

Recommended

Shows the recommended values for the following attributes of the VM:

- Logical CPUs
- Memory

- Memory Reservation
- CPU Reservation

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.
- **CPU Contention** - 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.
Host Name	Name of the host to which the VM belongs. Click the host name to navigate to the "Summary for Host" on page 51.
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 51.

Item	Description
Datacenter Name	Name of the datacenter to which the VM belongs. Click the datacenter name to navigate to the "Summary for Datacenter" on page 50 .
vCenter Name	Name of the vCenter to which the VM belongs.
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.
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(s)	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

To view the summary data for a particular datastore, from the Treemap, right click the datastore and select **Capacity Usage Trends**. The Summary page opens.

Alternatively, you can also access the page from the Optimization **Overview** page. Click the **DataStores** tab from the table at the bottom of the page. The page refreshes to display the Summary for Datastore page.

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)
- DB Status for VMs

Remarks

Provides information on the overall space and health of the datastore.

Domain Information

The table lists the following information:

- Virtualization Domain
- Datastore Type
- Hosts
- VMs
- I/O Size Analysis
- Disk Read Latency
- Disk Write Latency

Resource Information Table

The table lists the following information:

- 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 value is an aggregate of the disk usage of all the VMs corresponding to the resource.

- Disk Write Latency
- Disk Read Latency

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 ["Filtering Data" on page 48](#)

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 on the VM name to navigate to the ["Summary for VM" on page 54](#).
- **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.

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 ["Filtering Data" on page 48](#)

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 on the VM name to navigate to the ["Summary for VM" on page 54](#).
- **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.

CPU and Memory Reservation

This pages 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 ["Filtering Data" on page 48](#).

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 on the VM name to navigate to the ["Summary for VM" on page 54](#).
- **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.

Using Placement

You can use the Placement feature and 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.

To view the recommendation on where to allocate new VMs in your environment, follow these steps:

1. Click  from the **Options** pane in the vPV 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 **Reference VM CPU Size (GHz)** field.
4. Type the Memory size of the VMs in the **Reference VM Memory Size (GB)** field.
5. Type the Disk size of the VMs in the **Reference VM Disk Size (GB)** field.
6. Select the virtualization domain from the **Domain** field.
7. Click **Add VMs**.

The VM information appears in the **VMs list for placement** box.

8. Click **Find Placement**.

The **Placement Recommendation** for the VMs appears on the page.

Placement Recommendation

After providing the details of the VMs that you intend to add to your environment, vPV provides you the placement recommendation for the new VMs.

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 on the datacenter name, cluster name, or host name, you can navigate to the specific Summary page.

Example: 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 HP Virtualization Performance Viewer 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 Trends** option. The **Overview** page opens. Alternatively, you can click  from the vPV home page to 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	nmcvc.ind.hp.com
Clusters	2
Hosts	12
Powered On VMs	102
DataStores	38

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 46](#).

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 (NMC_490G&-A)

Cluster Name	CPU Available for Allocation(GHz)	CPU Usage(GHz)	CPU Reclaimable (GHz)
NMC-490G7-B	443.46	44.63	210.74
NMC-490G7-A	7006.11	35.63	530.53

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 46.

You can drill down and check the undersized VM. HP 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 HP 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 the all VMs that are available in the Datacenters, Clusters, or Hosts.
- Drill down to the VM where HP 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.

Chapter 5: Forecast

Based on historical usage trends, vPV can forecast the resource utilization of entities in your environment. You can use this data to plan for future resource requirements and allocations.

Navigation

To access the feature, click  from the Options pane in the vPV home page. The forecast page opens up.

Alternatively, you can right click an entity from the Treemap, and select **Forecast**. The forecast page for the particular entity opens up.

The Forecast window opens showing the data for the resources. After installing vPV, based on the usage trends of resources, vPV starts forecasting the resource utilization.

Note: The minimum requirement to get accurate graphs for forecast is that the data should be collected atleast 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.

vPV can forecast the following:

- CPU Utilization
- Memory Utilization
- Storage Utilization

Note: Storage Utilization data is not available for KVM VMs and Hosts.

You can view and drill down to the forecast data for the following entities:

- Aggregated Datacenter
- Aggregated Cluster
- Host
- VMs

When the number of VMs for any Host or User is increasing over time, even if the resource usage of each of the VMs are flat, the resource usage trend for the Host or User shows an increasing trend. In this case, the aggregated graphs for the Host or User shows increasing trends and may show a smaller value for Days To Capacity for the Host or User compared to the value for the individual VMs shown on the page. This is applicable for other aggregate graphs for Datacenter and Cluster.

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.

The tables corresponding to each graphs list the following:

- **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 vPV 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 to 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.

Drill Down

You can drill down to the entities in the datacenter and view the forecast data corresponding to each of the entities. Click on 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.

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 cluster aggregated view that the **Days to Capacity** is 8 days, then check the capacity of the hosts in the specific cluster.

Hosts		DataStores		
Host Names	Days To Capacity	CPU Capacity(GB)	Memory Days To Capacity	Memory Capacity(GB)
nnmw231.ind.hp.com	20	27.19	5	142.33
nnmw232.ind.hp.com	> 90	27.19	> 90	142.33

Here, one of the host is having days to capacity as 5 days. Take necessary actions to mitigate the issue.

Conclusion: If the data in the aggregated level shows a specific value, check the data of the entities available to make the necessary changes.

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 vPV.
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 "[Chapter 4: Optimization and Placement Overview](#)" on page 45

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 would reach 72.33 GHz after 90 days.
- 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 would reach 67.08 GB after 90 days.
- 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.

Chapter 6: 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 CSA with vPV" on page 20](#).

Follow the steps to log on as CSA user:

1. Browse to the CSA login page.
2. Type the CSA user name in the **Username** field.
3. Type the password in the **Password** field.
4. Select **CSA** from the Context field.
5. Select the required organization from the **Organization** field.
6. Click **Login**.

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

- Monitor the aggregated resource utilization of the available VMs
- Overview about resource utilization of all the VMs.
- Evaluate the capacity utilization based on the current trends

You can also get the aggregated capacity utilization of a VM for 30 days, 60 days, or 90 days. For more information, see [Forecast](#).

Chapter 7: 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 vPV.

Real Time Guest OS drill down tool (HP 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 hrs.

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](#).

Launch the Tool

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.

Note: When you select a VM and navigate for advance troubleshooting and Real Time Guest OS Drilldown is not available, you can download the tool from the message available on the screen. Follow the instructions available. For more information, see HP Compute Sensor user guide.

User Interface Details of Guest OS DrillDown

The following table lists the features on the Guest OS Drilldown page and the corresponding functions.

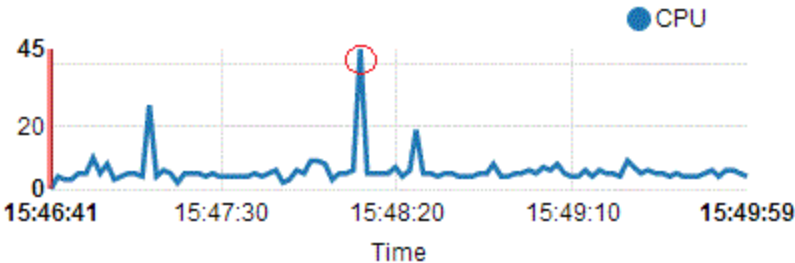
User Interface Details	Description
Graphs	Displays the graphs to check the resource utilization.
Tabs	Displays the details about the processes that are running, CPU usage, free disk space, swap memory, Filesystem, and Network usage running on VM.
Options	Contains the settings available on the Guest OS 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 Options.

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	<p>Shows the real time CPU utilization.</p>  <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.

Graph Name	Description
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 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 both utilization are high. • 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. You can get the updated



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 of the	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. On Windows, the process owner account name, without the domain name.

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 of the	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 process.
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 (KBs)	The data transferred to this disk device during the interval
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 System 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, space reserved.

Parameters of the File System tab

Name of the	Description
Device Number	Unique identifier of the device.
Device Name	Name of the process.
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 of the	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
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 of the	Description
Application ID	Unique identifier for the application.

Name of the	Description
Application Name	Name of the application.
Visible 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 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 System Events

Name of the	Description
Event ID	Unique identifier for the event.
Severity	The list for Windows is: Critical, Error, Warning, Information, or Verbose 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.
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.
Print	Click  to print the real time graphs.

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.

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 vPV Guest OS drill down tool to troubleshoot the reason for the same.

The steps include the following:

1. Log on to HP Virtualization Performance Viewer 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 **Processcloud** 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 **Processcloud** 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 8: Performance Graphing Overview

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.

Overview of Graphs

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.

Metrics

A metric is a measurement that gives an indication of the operational health and performance of a resource. Metric is a parameter or a set of parameters that you can use to monitor and measure the health, performance, and availability of a monitored resource.

Performance Graphing helps you to visualize the monitored data in a graphical or tabular format. A drawn graph consists of data points available for the selected metrics. A metric class is a set of related metrics grouped together based on the type of data the metric reports.

The metrics for KVM, Xen, OpenStack, and HP aPaaS are listed below along with their units. The metric units for VMware vCenter and HyperV are visible in the workbench.

KVM/Xen

Model Type	Metric Name	Description
Common Metrics <i>(These metrics are collected for both Hosts and VMs.)</i>	CPUTotalUtil	The total CPU utilization percentage. Note: The CPU Usage metric in the workbench is mapped to this metric.
	CPUPhysTotalTime	The total amount of time that a guest (VM) uses the host's CPU.
	CPUCycleTotalUsed	Total number of CPU cycles used.
	CPUClockSpeed	The CPU clock speed in MHz.
	MemSwapin	The swapped memory, in MB.
	MemPhys	The total physical memory in MB.
	NumCPUSocket	Number of CPU sockets.
	NumCPU	Number of CPUs.
Hosts	NumLS	Number of logical machines (guest VMs).
	MemEntlUtil	The percentage of memory utilization.
Virtual Machines	DiskPhysWriteByteRate	The rate at which data is transmitted, in KBps, to the physical disk.
	DiskPhysReadByteRate	The rate at which data is received, in KBps, from the physical disk.
	DiskPhysIOByteRate	The data transaction rate of the physical disk, in KBps.
	NetOutByte	Number of bytes, in MB, transmitted during the specified time interval.
	NetInByte	Number of bytes, in MB, received during the specified time interval.
	NetByteRate	Sum of data that is transmitted and received, in KBps.
	NumNetif	Number of network interfaces supported by a VM.
	NumDisk	Number of idle disks (including the CD ROM).

OpenStack

Model Type	Metric Name	Description
Tenant	NumVM	Number of VMs present in the tenant.
	NumHost	Number of Hosts present in the tenant.
	NumCPU	Number of CPUs present in the tenant.
	CPUUtil	The average CPU utilization rate, in percentage, of the VMs present in the tenant.
	DiskUtil	The average disk utilization rate, in percentage, of the VMs present in the tenant.
	MemoryUsageHours	<p>The amount of memory utilized by the VMs in the tenant during the hours that they were in an active state.</p> <p>For example:</p> <p>Consider a tenant which has 5 active VMs with 2 MB memory each. And the VMs are active for 10 hours.</p> <p>Then,</p> <p>MemoryUsageHours = (Memory capacity per VM) x (Number of VMs) x (Number of hours VMs are active) = 2 x 5 x 10 = 100.</p>
	CPUUsageHours	<p>The number of CPUs utilized by the VMs in the tenant during the hours that they were in an active state.</p> <p>For example:</p> <p>Consider a tenant which has 5 active VMs with 2 CPUs each. And the VMs are active for 10 hours.</p> <p>Then,</p> <p>CPUUsageHours = (No. of CPUs per VM) x (Number of VMs) x (Number of hours VMs are active) = 2 x 5 x 10 = 100.</p>

Model Type	Metric Name	Description
	DiskUsageHours	<p>The number of disks utilized by the VMs in the tenant during the hours that they were in an active state.</p> <p>For example:</p> <p>Consider a tenant which has 5 active VMs with 2 disks each. And the VMs are active for 10 hours.</p> <p>Then,</p> <p>DiskUsageHours = (No. of Disks per VM) x (Number of VMs) x (Number of hours VMs are active) = 2 x 5 x 10 = 100.</p>
	DiskUsageGB	The amount of disk space utilized, in GBs, by the VMs present in the tenant.
	InstanceLimit	The maximum number of VMs allowed in a tenant.
	CPULimit	The maximum number of CPUs allowed in a tenant.
	MemoryLimit	The maximum memory capacity allowed in a tenant.
	CPUUsageMHz	The total physical CPU utilization, in MHz, by the VMs present in the tenant.
	TotalCPUMHz	The total physical CPU capacity, in MHz, of the VMs present in the tenant.
	TotalDiskGB	The total disk capacity, in GBs, of the VMs present in the tenant.
Cloud	NumVM	Number of VMs present in the cloud.
	NumHost	Number of hosts present in the cloud.
	NumTenant	Number of tenants present in the cloud.
	CPUUtil	The CPU utilization percentage.
	MemUtil	The memory utilization percentage.
	NumCPUCores	The total number of CPU cores.
	Memory	The total memory capacity of the cloud, in MBs.

Model Type	Metric Name	Description
Virtual Machine	NumCPU	The number of CPUs allocated to the VM.
	NumDisk	The number of disks allocated to the VM.
	NumNIC	The number of NICs allocated to the VM.
	Memory	The total memory capacity of the VM, in MBs.
	DiskCapacity	The total disk capacity of the VM, in GBs.
	CPUUtil	The total CPU utilization percentage.
	MemUtil	The total memory utilization percentage.
	DiskUtil	The total disk capacity utilization percentage.
	NumDiskRead	The average number of disk read commands issued by the VM.
	NumDiskWrite	The average number of disk write commands issued by the VM.
	NumNetRead	The average number of network read commands issued by the VM.
	NumNetWrite	The average number of network write commands issued by the VM.
	CPUUsageMHz	The total physical CPU utilization, in MHz.
	MemUsageMB	The total physical memory utilization, in MBs.
	DiskUsageGB	The total disk utilization, in GBs.
	CPUPhysUtil	Total CPU utilization percentage of the physical host.
MemPhysUtil	Total memory utilization percentage of the physical host.	

Model Type	Metric Name	Description
Host	NumVM	The number of VMs hosted by the host.
	NumCPU	The number of CPU sockets available in the host.
	NumCPUCore	The number of CPU cores used by the host.
	NumDisk	The number of disk sockets available in the host.
	NumNIC	The number of NIC sockets available in the host.
	Memory	The total memory capacity of the host, in MBs.
	CPUUtil	The CPU utilization percentage.
	MemUtil	The memory utilization percentage.
	NumDiskRead	The average number of disk read commands issued by the host.
	NumDiskWrite	The average number of disk write commands issued by the host.
	NumNetRead	The average number of network read commands issued by the host.
	NumNetWrite	The average number of network write commands issued by the host.

HP aPaaS

Model Type	Metric Name	Description
Cloud	num_apps	The number of applications present in the cloud.
	num_users	The number of users present in the cloud.
	memory_totalGB	The total memory allocated to all the nodes in the cloud, in GBs.
	memory_free_totalGB	The total unutilized memory by all the nodes in the cloud, in GBs.
	memory_used_totalGB	The total memory utilization by all the nodes in the cloud, in GBs.
	memory_usage_pct	The total memory utilization percentage.
	cpu_avg_usage_pct	The average CPU utilization percentage.
	total_num_cpu	The total number of CPUs allocated to all the nodes in the cloud.
Group	name	The group name.
	members	The names of the users present in the group.
	user_count	The total number of users present in the group.
	memory_usage_pct	The total memory utilization percentage of all users in the group.
User	isAdmin	Information on whether the user is an administrator.
	emailId	The email address of the user.
	groupName	The group to which the user belongs.
	app_count	The number of applications used by the user.
	memory_allocated_mb	The total memory capacity allocated to the user, in MBs.
	memory_used_mb	The total memory utilization, in MBs.
	memory_usage_pct	The total memory utilization percentage.

Model Type	Metric Name	Description
Node	cpu_idle	The percentage of time that the CPU is idle.
	cpu_interrupt	The percentage of time that the CPU is interrupted.
	cpu_nice	The percentage of time that CPU uses the nice process.
	cpu_softirq	The percentage of time CPU encountered software interrupt requests.
	cpu_steal	The percentage of cycles that were 'stolen' from the CPU.
	cpu_system	The percentage of time the CPU is under system state.
	cpu_user	The percentage of time the CPU is under user state.
	cpu_wait	The percentage of time that the CPU is in wait state.
	df_freeGB	The free disk space, in GBs.
	df_usedGB	The used disk space, in GBs.













Model Type	Metric Name	Description
	fork_rate	The rate at which the fork system call is executed.
	net_rate	The throughput of the node, in bytes/seconds.
	net_out_rate	The number of bytes transmitted per second.
	net_in_rate	The number of bytes received per second.
	net_in_byte	The total number of bytes received by the node.
	net_out_byte	The total number of bytes transmitted by the node.
	net_out_error_pct	The percentage of packets which were not transmitted.
	net_in_error_pct	The percentage of packets which were not received.
	net_error_packets	The total number of packets which displayed errors during transmission and reception.
	net_in_error_packets	The total number of packets which displayed errors during reception.
	net_out_error_packets	The total number of packets which displayed errors during transmission.
	net_out_packets	The total number of packets transmitted by the node.
	net_in_packets	The total number of packets received by the node.
	num_netifs	The total number of network interfaces present in a node.

Model Type	Metric Name	Description
	load_longterm	The load average value during the last 15 minutes.
	load_midterm	The load average value during the last five minutes.
	load_shortterm	The load average value during the last one minute.
	memory_free_buff_cacheGB	The maximum memory buffer limit, in GBs.
	memory_usedGB	The memory utilization, in GBs.
	memory_usage_pct	The memory utilization percentage.
	ps_state_blocked	The total number of processes in blocked state.
	ps_state_paging	The total number of processes in paging state.
	ps_state_running	The total number of processes running on the node.
	ps_state_sleeping	The total number of processes in sleep state.
	ps_state_stopped	The total number of processes which are stopped.
	ps_state_zombies	The total number of processes in the zombie state.
	swap_cachedGB	The total size of the swap cache, in GBs.
	swap_freeGB	The free swap space , in GBs.
	swap_io_inGB	The amount of data swapped from the disk, in GBs.
	swap_io_outGB	The amount of data swapped to the disk, in GBs.
	swap_usedGB	The amount of swap memory used, in GBs.
	num_cpu	The total number of CPUs allocated to the node.

Model Type	Metric Name	Description
Apps	name	The name of the application.
	environment	The application environment variable details.
	instances	The total number of application instances.
	runningInstances	The total number of application instances which are running.
	services	The list of available and provisioned service names.
	state	The state of the application (Started/Stopped).
	uris	The unique ID of the application.
	version	The version number of the application.
	stats_state	The status of the application instance.
	cores	The number of CPU cores allocated to the application.
	disk_quote_gb	The disk capacity allocated to the application, in GBs.
	host	The host on which the application is deployed.
	port	The port number of the application.
	mem_quota_mb	The total memory capacity allocated to the application, in MBs.
	uptime	The total time for which the application is running.
	cpu_usage_pct	The CPU utilization percentage.
	mem_usage_mb	The memory utilization, in MBs.
	mem_usage_pct	The memory utilization percentage.
	disk_usage_gb	The disk utilization, in GBs.
	disk_usage_pct	The disk utilization percentage.
	time_usage	The amount of time for which the application is in use.

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
	<ul style="list-style-type: none"> • Host (VMware vCenter, HyperV, KVM, Xen, and OpenStack) • Node (HP aPaaS)
	Host in maintenance mode
	Storage
	Host Group
	<ul style="list-style-type: none"> • Cloud (OpenStack) • Target (HP aPaaS)
	<ul style="list-style-type: none"> • Tenants (OpenStack) • User Groups (HP aPaaS)
	User
	Application







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

Workbench Overview

The Workbench page provides the Performance Graphing functionality for vPV.



Configuration Items

The Configuration Items tree displays the available 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. 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 104](#).

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	<p>Displays the following information:</p> <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 <p>You can use the Refresh () icon available in the Metrics tab to refresh the available metric classes, instances (if any), and metrics.</p>
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 106




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 and Metrics tabs.
Expand (>>)	Click to show the Graphs and Metrics tabs. This icon appears after you collapse the tabs.
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 draws graphs on the Performance pane.
Save as Favorite ()	Click to save drawn graphs as a Favorite and access later. This icon appears after you draws 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 . vPV 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 are 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 " on page 101.
Options > Navigate	Select this option to select the start and end date from the Time Settings dialog box. The graph data 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 100
 (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 is refreshes at a specific interval.
 (Prev/Next)	Performance Graphing provides 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.

- 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.

- 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:

- Click () **Table Highlights** icon. The Table Highlights dialog box opens.
- Select the required value based on which the cells must be highlighted, from the drop down list.
- Select the required comparison symbol from the list of symbols available. For more information on the comparison symbols, see "[Comparison Symbols](#)" below
- Type the value to be compared in the text box.
- Select a color from the color palette to highlight the cells.
- 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**.

- 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)

Comparison Symbol	Description
=	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 opens.
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:

- If you are selecting a Microsoft Excel or a .tsv graph, you must configure the browser settings to display Microsoft Excel and .tsv files. For more information on browser settings, see "[Chapter 9: Troubleshooting vPV](#)" on page 109.
- The exported data has the same time as that of the vPV 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)	<p>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 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>


, continued

Icon/Option	Description
 (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
All	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** . vPV 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](#)" below.

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

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 new window)	Opens the report, in the specified HTML or PDF format, in a new 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

vPV 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.

Chapter 9: Troubleshooting vPV

The following section details how to troubleshoot vPV:

Unable to scroll to bottom of browser window

Symptom	When accessing vPV 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 vPV interface is 1280x768 and the browser must be in the full screen mode.

Unable to log on to the vPV interface after logging out

Symptom	When you log out and try to log on to the vPV user interface again, the page refreshes to display the Login page again.
Resolution	Provide your credentials and try to log on again. The vPV 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. 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 vPV, you are not able to access vPV from VMware vCenter Server.
Cause	There may be a firewall between vPV server and VMware vCenter Server.
Resolution	Configure the firewall to allow communication between vPV server and the VMware vCenter Server. The ports that must be opened are 80 and 443.

Last Collection Time in Administration page does not match with the last data time in Treemap and graphs

Symptom	The last collection time displayed on the Administration page is different from the last data time shown in the treemap and graphs.
Cause	When you are using vPV to monitor Microsoft SCVMM data sources, the last collection time in the Administration page reflects the time from the vPV server and the graph data points show the SCVMM server time. If both the servers are in the same time zone but the time is not synched, this problem occurs.
Resolution	When the vPV and Microsoft SCVMM servers are in the same time zone, make sure to synch the time on the servers.

Unable to connect to Virtual Machine Manager Management server local host


Symptom	The <code>getscvmpdata.log</code> file shows the error message <i>Unable to connect to VMM management server localhost</i> .
Cause	The Virtual Machine Manager service on the server is not responding.
Resolution	Check if the SCVMM services are running. If not, start the SCVMM service.

Unable to access vPV in BSM when it is accessed using HTTPS

Symptom	When vPV 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? '. Upon clicking 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	To disable the mixed content warning: <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. 5. Click the Disable radio button. 6. Click OK. <p>The mixed content warnings are disabled.</p>

Collection is not working in KVM/Xen

Symptom	Data is not collected for one or more data sources.
Cause	<ul style="list-style-type: none"> • ssh is not automated • The trace logs in <code>ovpmtrace.0.txt</code> are not logged properly. • There are 2 opsagt processes running on the vPV 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. <p>To check whether ssh is automated:</p> <ol style="list-style-type: none"> a. Run the command: <pre>ssh <username>@<KVM/Xen hostname></pre> <p>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.</p> b. If you are prompted for a password, then the collection is not working. <ul style="list-style-type: none"> • Run the command: <pre>For KVM: virsh -c qemu+ssh://root@<hostname>/system</pre> <pre>For Xen: virsh -c xen+ssh://root@<hostname>/</pre> <p>Where, <i><hostname></i> is the host name of the KVM/Xen server.</p> <p>If you are prompted for a password, then the collection is not working.</p> • Check for trace logs in ovpmtrace.0.txt located at /var/opt/OV/log/. <p>Solution 1</p> <p>To restart the collection:</p> <ol style="list-style-type: none"> 1. Click  on the Options pane. The Administration page opens. 2. On the Data Sources tab, click Restart Collection. <p>Solution 2</p> <p>If there are 2 pvcd processes running on the vPV 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: <pre>ovc -stop pvcd</pre> 2. It is recommended that you update the libvirt library to the latest version on the vPV machine. 3. Restart collection using the following command: <pre>ovc -start pvcore</pre>
------------	--

Unable to launch Real Time Guest OS Drilldown, port is not available

Symptom	When accessing HP vPV 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 Windows node - %HPCSInstallDir% ■ On Linux or Linux Debian node - /opt/ OV/ hpcs 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.

Collection is not working in HP aPaaS

Symptom	Data fails for HP aPaaS data sources.
Cause	HP aPaaS target is not accessible from the vPV host.

Resolution	<p>To check whether HP aPaaS target is accessible:</p> <ul style="list-style-type: none"> • Try pinging the HP aPaaS target from the vPV host: <pre>ping api.abcd-azd7.local</pre> <p>If there is no response from the target, then the collection is not working.</p> <p>To enable collection from the host:</p> <ol style="list-style-type: none"> 1. Open the hosts file from <code>\etc\</code> directory. 2. Add the appropriate host information in the hosts file. <pre><host IP address><HP aPaaS target URL></pre> <p>where,</p> <p><code><host IP address></code> is the host IP address and,</p> <p><code><HP aPaaS target URL></code> is the host URL.</p> <p>For example:</p> <pre>12.345.67.890 api.abcd-azd7.local</pre> <ol style="list-style-type: none"> 3. Ping the HP aPaaS target again to see if connection is working.
------------	---

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	vPV server time and vCenter time may not be in sync.
Resolution	Check if the vPV server time and data source time are in sync. vPV collects time bound data from the vCenter. So time must be synchronized across the vPV virtual appliance and the vCenter appliance or server.

Not able to access vPV user interface when PKI authentication is enabled

Symptom	When trying to access vPV 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 vPV 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 vPV using Google Chrome.

Graphs show less than 12 data points (based on 5 minute data collection) for an hour

Symptom	<p>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> <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"> vPV 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 vPV 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 vPV. The huge volume of data, might result in the delay in data logging due to indexing and co-relation between entities.
Resolution	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 vPV.

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	If the vPV Virtual Appliance is shutdown abruptly, sometimes the database does not start. Due to this, <i>pvcd</i> exits on rebooting.
Resolution	Restart the <i>pv</i> process (<i>pvcd</i> , Tomcat, and Vertica) after rebooting the machine. Run the command, pv restart from the vPV console.

Data collection is not getting completed

Symptom	In the Data Source tab of Administration page, data collection status for vCenter shows as Data collection in progress for long time.
Cause	If the Cluster names are not unique across the vCenters monitored by vPV, data collection does not get completed.
Resolution	Ensure that the Cluster names are unique across your vCenters.

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 /opt/OV/support/ovtrccfg -cf /opt/OV/newconfig/conf/xpl/trc/pvcdXPL.tcf The XPL trace files are present at /tmp/pvcd_*.trc.</p> <p>To stop tracing, run /opt/OV/support/ovtrccfg -app pvcd -off</p>	<p>If tracing is enabled, /tmp/ directory contains the trace files (use - ls /tmp/pvcd_*.trc)</p>
vPV 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 vPV, when data collection does not happen properly, the system displays some error messages. The following table lists the error messages, scenarios when these messages can appear, and resolution, wherever applicable.

Error Message /Status Message	Scenario	Resolution
Connecting...	vPV is establishing connection with the data source.	No resolution required.
Connected	Collection is completed successfully and waiting for the next collection to start.	No resolution required.
Data Collection in Progress...	Collecting data and storing it in the database.	No resolution required.

Error Message /Status Message	Scenario	Resolution
<p>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 Administration page.</p>	<p>Instance count exceeds the specified limit of the free license.</p>	<p>Perform one of the following actions:</p> <ul style="list-style-type: none"> • Wait for the next data collection to complete. • Remove a data source to adjust the instance count. For more information on removing data sources, see "Removing a Data Source" on page 18. • Install a permanent license with sufficient capacity or enable the evaluation license from the Administration page.
<p>Total number of instances exceed the maximum limit of the license instance count (<maximum count>). Use a data source or data source credentials with fewer instances associated.</p>	<p>Instance count exceeds the specified limit of the free license.</p>	<p>Perform one of the following actions:</p> <ul style="list-style-type: none"> • Use a data source or data source credentials with fewer instances associated. • Remove a data source to adjust the instance count. For more information on removing data sources, see "Removing a Data Source" on page 18.
<p>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.</p>	<p>Instance count exceeds the specified limit of the free license.</p>	<p>Perform one of the following actions:</p> <ul style="list-style-type: none"> • Wait for the next data collection complete. • Restart using the 'Restart Collection' button on the Administration page.

Error Message /Status Message	Scenario	Resolution
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 data source to adjust the instance count. For more information on removing data sources, see "Removing a Data Source" on page 18. • Limit the number of instances associated with a data source.
Data Collection failed. Check if Libvirt services are running on the vPV server.	The Libvirt services are not running on the vPV server.	The Libvirt services must be running on the vPV server.
Invalid username or password	Credentials provided are wrong.	Provide the correct credentials.
Connection Timeout	Connection request to the data source has timed-out.	Wait for the next data collection.
Unable to connect to vCenter server	This can be due to any of the following scenarios: <ul style="list-style-type: none"> • Unable to connect to data source • Invalid data source name • System is not on the network. 	Actions: <ul style="list-style-type: none"> • Check if the data source is running. • Provide the correct data source name. • Ensure that the data source is accessible from vPV.

Error Message /Status Message	Scenario	Resolution
Unable to collect performance data from the vCenter	No performance data is coming from the data source.	Make sure that the data source services are up and running and historical data collection is happening.
Unable to connect to the Database	Database connection fails.	Try restarting vPV.
Unable to start collection: Error occurred while connecting to the Database	Database is not running or vPV is not able to connect to the database.	Try restarting vPV.
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 while reading the database configuration file.	Contact Support.
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> .

Chapter 10: Frequently Asked Questions

All Data Sources

- Why do some graphs show a message "No data found for the duration with points every 5 minutes...." when I click on 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 vPV utilize too much CPU to run?
By default, vPV collects data from the data sources every 5 minutes. If a data source response is slow or vPV 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.

No data is available for the hosts, host groups, and host clusters in the Microsoft SCVMM domain from the following metrics:

- StorageIOPSUsage_bytesPerSecond
- NetworkIOUsage_bytesPerSecond
- Why am I not able to see folders in **vPV**?
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 data sources of the same type contains instances with same name, then the vPV 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 **vPV**?

To check the status of the database:

- a. Run the following commands on the vPV 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 data source displayed in the treemap and the instance count of the same data source displayed on the Licensing and Data Management tabs in the Administration page.

When a data source goes in an error state, the instances from that data source are not considered while calculating the instance count displayed in the Administration page.

However, the instances from the erroneous data source will still be considered while calculating the instance count displayed in the vPV treemap as long as the data source is not deleted from the vPV Administration 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 data sources added in vPV.

- The **vPV** treemap shows data for a data source even if the data source is in an error state. The vPV treemap shows the data for a data source in error state for one hour after the last collection time.
- The vPV title bar hides the performance graphs' title bar if the performance graphs are dragged to the vPV title bar. This prevents me 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 vPV?

If you do not see data for datastores, wait for one hour and then check the data again.

- What are the log files available with vPV?

The following table lists the log files available with vPV:

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
vPV Web Server	vPV 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?
 - Open the file
OVINSTALLDIR/newconfig/OVPM/smepack/VCENTER/integration/VCENTER_GC_Integration.xml.
 - Locate the corresponding **CI_VIEW** tag (for example, VM (Compute)) and **METRIC** tab (for example, guestMemPercent) for which you want to set the thresholds.
 - 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 properly. 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 vPV, 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 **Data Sources** tab of the **Administration** page) to see the resources of the VMware vCenter Server in the Treemap.

- Why does not Datastore metric appear in vPV?

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 data stores?

You may not see micro charts for data stores 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.

- 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 data stores 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 data stores.
- Why do the default graphs for data stores 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 **vPV**?
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 **vPV**?
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 Data Source tab, the Data Sources tab 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.

Microsoft SCVMM

- Why does vPV draw performance graphs for Powered Off VMs?
Sometimes, when a VM is deleted from vCenter, the sync between SCVMM and vCenter does not happen properly. So the performance graph for the powered off VM is still drawn.
- Why do I see the error - Error occurred: Unable to connect to remote server, in the log file?
There might be a firewall between vPV server and Microsoft SCVMM. Configure the firewall to enable the communication and the ports enabled must be 8081 and 8444.



- I have removed the Microsoft SCVMM data sources from the Administration page. I still see data in the Treemap.
This is because you need to run the **Stop-Collector.bat** file before removing the data source. For more information, see "[Removing a Data Source](#)" on page 18
- Why do I not see micro charts for VMs?
The micro charts are drawn only after two data collection cycles are completed.
- I deleted a VM from Microsoft SCVMM but I still see the VM in the Configuration Items tree in the Workbench page.
After deleting a VM or host from SCVMM, you can still see the VM or host in the Configuration Items tree in the Workbench. But it is grayed-out and indicates that currently the resource is in the powered off state or data is not collected from the resource. If you draw graphs or launch reports against such a resource, data shown is till the time data was collected from the resource.
- Data is not visible for SCVMM data source immediately after the data source are added.
Sometimes the data collection interval for SCVMM is more than 9 minutes.

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.
vPV 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 data source display CPU utilization as 100%
The guests will show the correct CPU utilization after 2 collection cycles.

HP aPaaS

- The  icon for Users displays incorrect details about the Users.
In the HP aPaaS domain, if there are no applications available, then the  icon for Users in the Treemap displays incorrect details about the Users.

- Despite adding multiple HP aPaaS targets to vPV, the treemap displays only one node.
vPV takes at least 5 minutes to display all the nodes if you add multiple HP aPaaS targets.
- When multiple HP aPaaS hosts are deleted, the Data Sources tab still displays a few of the deleted hosts.

Close the Internet browser and open it again.

We appreciate your 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 (Virtualization Performance Viewer 1.20)

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 your_IE_team_PDL@hp.com.