

Power Systems

IBM PowerVP

IBM

Power Systems

IBM PowerVP

IBM

Note

Before using this information and the product it supports, read the information in "Notices" on page 29.

This edition applies to IBM AIX Version 7.1, to IBM AIX Version 6.1, to IBM i 7.2 (product number 5770-SS1), to IBM Virtual I/O Server Version 2.2.4.00, and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2014, 2015.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Getting started with PowerVP	1
What's new in IBM PowerVP	1
PowerVP prerequisites	2
PowerVP system considerations	2
PowerVP installation	3
Installing PowerVP	3
Installing the AIX and VIOS agent	6
Installing the PowerVP GUI application	7
Installing the IBM i agent	7
Installing the Linux agent	7
Additional instructions for Linux	9
Configuring SSL on PowerVP.	9
Configuring SSL on AIX, VIOS, or Linux	9
Configuring SSL on IBM i	10
Configuring SSL on partition-level agents	10
Starting and using the PowerVP GUI.	11
Configuring the PowerVP agent	11
Starting the PowerVP GUI	11
Connecting to the agent	11
Thresholds and alerts	12
Permanent data store	13
Configuration file directives	14
Data store format	14
Timestamp entries	15
PowerVP Information entries	15
Topology	16
System topology.	16
Chip topology	17
Core topology	17
Registered partitions	17
Affinity information by domain	18
Affinity information by partition	18
Affinity domain information by virtual processor.	18
System Agent Data	18
CPU Utilization by core	18
Timebase cycles by core	19
Bus utilization by Timebase cycles by core	19
CPU utilization by partition	19
Partition metrics.	19
Partition Agent Data	19
Partition status	20
Partition cycles per instruction (CPI)	20
Partition Ethernet throughput	20
Partition disk I/O	20
Partition CPU cycles	20
Recording and playing performance metrics in PowerVP	20
Recording performance metrics.	21
Replaying recorded performance metrics	21
Using PowerVP to access the VIOS Performance Advisor tool	22
PowerVP security considerations	22
Using PowerVP to investigate and optimize performance	23
General terms used with PowerVP	23
PowerVP FAQ	25

Notices	29
Privacy policy considerations	31
Programming interface information	31
Trademarks	31
Terms and conditions	31

Getting started with PowerVP

Power[®] Virtualization Performance (PowerVP[™]) is a performance monitoring solution that provides detailed and real-time information about virtualized workloads that are running on IBM[®] Power Systems[™] servers. PowerVP is a licensed program that is offered as part of PowerVM[®] Enterprise Edition, but is also available separately for clients without PowerVM Enterprise Edition. You can use PowerVP to understand how virtual workloads use resources, to analyze performance bottlenecks, and to make informed choices about resource allocation and virtualized machine placement.

The PowerVP tool monitors the performance of an entire system (or frame). PowerVP is supported on AIX[®], IBM i, Linux, or Virtual I/O Server operating systems. It provides a graphical user interface that you can use to monitor virtualized workloads. PowerVP includes a system-level monitoring agent that collects data from the PowerVM hypervisor, which provides a complete view of virtualized machines that are running on the server. PowerVP displays the data that is collected at the system level, at the hardware node level, and at the partition level. You can optimize performance by using the PowerVP performance metrics, which provide information about balancing and improving affinity and application efficiency.

PowerVP provides an illustration of the Power Systems[™] hardware topology along with resource usage metrics. The resources utilization is portrayed using different colors to represent utilization thresholds. For example, green color indicates normal, yellow color indicates caution, and red color indicates an action or resource adjustment should be considered. The colors and thresholds are customizable because the thresholds are installation-dependent. Metrics include nodes, processor modules, cores, power bus links, memory controller links, GX I/O bus details, disk drives, and Ethernet links.

The PowerVP tool provides a mapping between real and virtual processor resources. The PowerVP agent, running on your Power system can be configured to "record" the PowerVP performance data in a file. This information can then be played by using the PowerVP graphical user interface digital video recorder (DVR)-like functions such as play, fast forward, rewind, jump, pause, or stop. You can find performance bottlenecks by replaying the recorded data at any point in time.

What's new in IBM PowerVP

Read about new or significantly changed information in IBM PowerVP since the previous update of this topic collection.

October 2015

The following updates have been made to the content:

- Starting with version 1.1.3, the PowerVP Graphical User Interface (GUI) is a web-based application that is viewed with a browser. The standalone Java PowerVP GUI is not available in PowerVP version 1.1.3 or later.

June 2015

The following updates have been made to the content:

- Added new information about using thresholds and alerts in PowerVP in "Thresholds and alerts" on page 12.
- Added a new section about using permanent data store to log PowerVP performance data to a file in "Permanent data store" on page 13.
- Added new information on using PowerVP to access the VIOS Performance Advisor in "Using PowerVP to access the VIOS Performance Advisor tool" on page 22.

PowerVP prerequisites

PowerVP runs on the IBM POWER7® and IBM POWER8® processor-based servers. On POWER7 PowerVP requires the firmware release to be at version 7.7 or later, depending on the POWER7 system model. On POWER8, PowerVP runs on all releases of the firmware. However, initial firmware levels cannot report the POWER® bus utilization. Therefore, the busses do not display any utilization. Later levels of the POWER8 firmware display the POWER bus utilization. The PowerVP agent runs on AIX, VIOS, IBM i, and Linux.

PowerVP supports IBM AIX version 6.1 and AIX version 7.1. To view the cycles per instruction (cpi) information on a POWER8 system, ensure that the correct service pack for the AIX operating system is installed.

PowerVP supports VIOS versions that are included with the POWER7+™ and POWER8 processor-based servers. To view the cpi information on a POWER8 system, ensure that the correct version of VIOS is installed.

PowerVP supports RedHat Enterprise Linux (RHEL) 6.4, or later, SUSE Linux Enterprise Server 11 SP3, or later and Ubuntu Linux15.04. On POWER8 processor-based servers, the partition drill-down information, such as the cycles per instruction (cpi) and the detailed view of the load and store unit cycles per instruction (LSU cpi), is not available on currently available Linux distributions.

PowerVP support IBM i version 7.1 (with technology refresh 8 (TR8) and program temporary fix (PTF) SI52700) and IBM i version 7.2 (with PTF SI53120). The PTFs have prerequisite and corequisite PTFs that are included when you order them. All requisite PTFs must also be applied to the IBM i server.

The Graphical User Interface (GUI) client is a web-based application that runs in an application server and renders to your web browser. PowerVP bundles a version of WebSphere® Application Server Liberty, which can be installed on a client workstation and the PowerVP application is deployed to the installed Liberty server. You can also install the PowerVP application to an application server of your choice, which allows you to make the PowerVP GUI available in your network from your own web pages. Multiple browsers can connect and view data simultaneously; however, each connected browser adds workload to your system because it requests performance data from the agent every second. The PowerVP agents collect data only when a GUI client is connected to it or when you are recording the performance data in the agent.

The GUI client requires Java™ version 1.7 on the application server and on the client computer with the browser. The GUI client is supported on the following browsers:

- Microsoft Internet Explorer, version 11 or later
- Mozilla Firefox, version 38 or later
- Google Chrome, version 44 or later

PowerVP system considerations

PowerVP requires at least one of the logical partitions (LPARs) on your POWER processor-based server to be identified as a partition that will collect the system-wide information for all partitions. This is known as the system-level agent. The remaining partitions are known as partition-level agents.

Review the following PowerVP configuration examples.

1. If a system has two VIOS partitions, three AIX partitions, five IBM i partitions, and two Linux partitions, the VIOS partitions are redundant and provide virtualization to other partitions. For the VIOS partitions to be also redundant to PowerVP, you must load the PowerVP agent for AIX/VIOS on both the VIOS partitions and configure them as system-level agents. On the AIX partitions, you can load the PowerVP agent for AIX/VIOS, configure them as partition-level agents, and you can specify both of the VIOS partitions as the associated system-level agent. On the IBM i partitions, you

can load the PowerVP agent for IBM i, configure it as the partition-level agents, and you can specify both of the VIOS partitions as the associated system-level agent. On the Linux partition, you can load the PowerVP agent for Linux, configure it as a partition-level agent, and specify both of the VIOS partitions as the associated system-level agent. You can start the PowerVP GUI and connect the GUI to either of the VIOS system-level agents and you can view all of the partitions except the other VIOS partition. If you connect the GUI to both VIOS system-level agents, you will be retrieving the same data from both systems but it will not have been collected necessarily at the same time. Therefore, you might not view the exact same data when you look at the two PowerVP GUI screens. By default, the performance data is collected every second, but the timing of the collection in the PowerVP agent and in the PowerVP GUI will not be synchronized.

2. To view the partition drill down data on a system with one VIOS, two IBM i, two AIX, and two Linux partitions, you can load the PowerVP agent for AIX/VIOS on the VIOS partition and you can configure it as the system-level agent (you can configure one or two system-level agents). On the other partitions, you load the corresponding PowerVP agent and configure them as partition-level agents with the VIOS partition as the associated system-level agent.
3. If a system has two IBM i, two Linux, and one AIX partitions, and no VIOS partition, you can choose any of the 5 partitions as the system-level agent. The system-level agent need not be a VIOS partition. You would load the PowerVP agent on all of the partitions, and configure one or two of the partitions as system-level agents and configure the other partitions as partition-level agents.

The system-level agent must have the **Allow performance information collection** partition property set. You can set this partition property by using the Hardware Management Console (HMC). Ensure that the **Allow performance information collection** partition property is checked in the **General** tab.

The PowerVP agent can be installed on any partition(s) of your POWER7 or POWER8 processor-based server. The partitions that you identify as the system-level agent(s) should be installed first. If you want to view partition-specific information for a partition, you must install and run the PowerVP agent on those specific partitions, which will be referred to as partition-level agents. The system-level agent also acts as a partition-level agent. The partition-level agents are configured to point to the system-level agent by using the TCP/IP host name of the system-level agent partition. The partition-level agents must connect to the system-level agent. Therefore, the system-level agent must be running before the partition-level agents can collect and provide partition-specific information. For the PowerVP GUI to display information about the system and its partitions, the system-level agent must also be running. You can start the partition-level agents before starting the system-level agent. The partition-level agent will retry to connect to the system level agent.

PowerVP installation

You can install PowerVP by using an installation wizard. The PowerVP GUI application can be installed along with IBM WebSphere Application Server Liberty on your client, or you can choose to have the installer only copy the PowerVP application to the client and you can deploy it to the application server of your choice. The PowerVP installer creates the selected shortcuts or menu options on the client where you run the installer. The PowerVP agent for IBM i is automatically installed on IBM i partitions by using the Restore Licensed Program (RSTLICPGM) with remote commands. The PowerVP agent for AIX and VIOS is packaged in the installp format. The installp file set must be copied to the AIX partition to complete the installation. The PowerVP agent for Linux is packaged in the RPM Package Manager (RPM) format. The RPM packages must be copied to the Linux partition to complete the installation. You can use File Transfer Protocol (FTP) to move the files to the AIX, VIOS, or Linux partition. You must transfer the files in binary mode.

Installing PowerVP

Complete the following steps to install PowerVP.

Insert the installation media into a machine that is running the Windows, Linux, or AIX. On the installation media, change the current directory to the directory of your client. The PowerVP and

PowerVP_IBMi_Agent executable files are in that directory. You must run the PowerVP executable file to complete the installation. You must run the executable file from the installation media.

1. Run the PowerVP executable program.
2. In the Installation window, select the language to be used during installation and click **OK**.
3. In the Introduction window, review the introduction information and click **Next**.
4. In the Software License Agreement window, read the license agreement and click **Accept**, and then click **Next**.
5. In the Choose PowerVP Install Set window, select one of the following options and click **Next**:
 - **Typical**
 - **PowerVP Client GUI**
 - **PowerVP Server Agents**
6. In the Choose PowerVP Install Set window, select the destination folder for PowerVP and click **Next**.
7. In the Choose PowerVP Shortcut Folder window, select the location for the PowerVP product icons and click **Next**.
8. Select whether you want the PowerVP installer to install a bundled copy of IBM WebSphere Application Server - Liberty to your workstation and deploy the PowerVP application to this application server. You can also choose to have the installer copy the PowerVP GUI application file to your client so you can deploy it to your own application server.
9. In the PowerVP GUI Configuration Summary window, review the summary information and click **next**.
10. In the PowerVP Agent Server Selection window, select the operating systems in which you want to install the PowerVP Agent and click **Next**.
11. Depending on your operating system selection in the PowerVP Agent Server Selection window, choose one of the following options:
 - For AIX/VIOS PowerVP agent, go to step 12.
 - For IBM i PowerVP agent, go to step 15 on page 5.
 - For Linux PowerVP agent, go to step 19 on page 5.
12. If you selected AIX/VIOS, the PowerVP AIX/VIOS Agent Specification window asks for the information that is needed for the AIX/VIOS installation.
 - The first field, **System Level Agent Hostname or IP Address**, is the TCP/IP hostname of the system level agent partition.
 - The **Default Agent Port** and **Default SSL Agent Port** are the default TCP/IP ports that are used by the GUI and agents to exchange performance data. If the defaults of 13000 and 13001 are used by another application in your installation, you can provide different ports that use these 2 fields.
 - The **Use SSL** and **Use non-SSL** check boxes are used to indicate which type of connection you want between the GUI and the agents. The PowerVP product can be configured to use either one or both depending on the needs of your environment. If you select the **Use SSL** check box, you need to complete the SSL configuration on your AIX/VIOS partition after it is installed. This configuration includes providing a digital certificate for the server and selecting the other SSL options for SSL connections. See “Configuring SSL on PowerVP” on page 9 for more details on SSL configuration.
 - The **Autostart Agent Server** check box indicates whether you want the PowerVP agent to be automatically started on the partition you are installing it to. On AIX and VIOS, PowerVP adds a startup script, `SPowerVP`, to the `/etc/rc.d/rc2.d` directory to automatically start the PowerVP agent at IPL.
 - The **Launch Installation Instructions** check box, when checked, launches a README dialog box in your web browser with information about installing the AIX and VIOS agent. To complete the AIX/VIOS installation, you use `installp`.
13. Click **Next** when you finish providing information in the PowerVP AIX/VIOS Agent Specification window.

14. In the PowerVP AIX/VIOS Agent Configuration Summary window, review the summary information and click **Next**.
15. The IBM i Agent System Information window is used to provide IBM i system information. For IBM i, the installation program automatically performs the licensed program installation on the IBM i partition. You can specify multiple IBM i partitions to be installed simultaneously. Click **Add system**.
16. The IBM i Agent System window will appear that contains prompts for agent installation information.
 - The first field, **IBM i System Hostname or IP Address**, is the hostname or IP address of the IBM i partition that you are installing the agent on.
 - The next field, **System Level Agent System**, is the TCP/IP hostname or IP address of the system level agent partition. If you are installing the agent on a partition that will be a partition level agent, you need to specify the hostname of the system level agent partition. If this partition is the system level agent, you use the same name that you provided in the **IBM i System Hostname or IP Address field**.
 - The **Default Agent Port and Default SSL Agent Port** are the default TCP/IP ports that are used by the GUI and agents to exchange performance data. If the defaults of 13000 and 13001 are used by another application in your installation, you can provide different ports that use these 2 fields.
 - Your IBM i partition user ID and password are required for the remote installation (the user name must have authority to the Restore Licensed Program (RSTLICPGM) command or have *SECADM or *ALLOBJ authority).
 - To use SSL, select the **Use SSL** check box. If you select the **Use SSL** check box, you need to complete the SSL configuration on your IBM i partition after it is installed. This configuration includes providing a digital certificate for the server and selecting the other SSL options for SSL connections. For more information about this, see the PowerVP SSL Configuration section in “Configuring SSL on PowerVP” on page 9.
 - The **Autostart Server Agent** check box indicates whether you want the PowerVP agent to be automatically started on the partition you are installing it to. On IBM i, PowerVP is configured as an Autostart TCP/IP server. If you do not select to have it auto-started, you can use the Start TCP Server (STRTCPSVR) and End TCP Server (ENDTCPSVR) CL commands to start and end the PowerVP agent (the SERVER parameter value is *POWERVP).
 - The **Launch Installation Instructions** check box, when checked, launches a README dialog in your web browser with information about the IBM i agent.
17. Click **OK** when done with IBM i system agent. If you have more IBM i systems to install, you can click the **Add System** button to add them. When done adding systems, click **Next**.
18. Review the information on the IBM i Agent System Pre-Installation Summary window and click **Next**.
19. If you selected Linux, the PowerVP Linux Agent Specification window is used to provide Linux partition information.
 - The first field, **System Level Agent Hostname or IP Address**, is the TCP/IP hostname of the system level agent partition.
 - The **Default Agent Port and Default SSL Agent Port** are the defaults that are used by the GUI and agents to exchange performance data. If the defaults of 13000 and 13001 are used by another application in your installation, you can provide different ports that use these two fields.
 - The **Use SSL** and **Use non-SSL** check boxes are used to indicate what type of connection you want between the GUI and the agents. The PowerVP product can be configured to use either one or both depending on the needs of your environment. If you select the **Use SSL** check box, you need to complete the SSL configuration on your Linux partition after it is installed. This configuration includes providing a digital certificate for the agent and selecting the other SSL options for SSL connections. For more information about this, see the PowerVP SSL Configuration section of “Configuring SSL on PowerVP” on page 9.
 - The **Autostart Agent Server** check box indicates whether you want the PowerVP agent to be automatically started on the partition you are installing it to. On Linux, PowerVP adds an ibm

powervp service definition to the /etc/init.d file which can be set to automatically start the PowerVP agent at IPL. If you do not select to have it auto-started, you can use “service ibm_powervp start” to start the PowerVP agent.

- The **Launch Installation Instructions** check box, when checked, launches a README dialog in your web browser with information about installing the Linux agent. You use an RPM to complete the Linux installation.
20. Click **Next** when done with the PowerVP Linux Agent Specification window.
 21. The PowerVP Linux Agent Configuration Summary window shows a review panel for Linux with your selections. Click **Next** if everything is correct, or click **Previous** to make any changes.
 22. In the Pre-Installation Summary window, review the summary information and click **Install**.
 23. In the Install Complete window, click **Done** to exit the installer.
 24. If you selected the option to launch the installation instructions, the instructions for completing the PowerVP agent installation on AIX/VIOS, instructions to complete PowerVP agent installation on Linux, and additional information for installation of PowerVP agent installation on IBM i, are displayed in your default web browser.

Installing the AIX and VIOS agent

After the GUI installation is complete, there are more steps that are required to complete the installation on the AIX and VIOS agent.

To install the AIX and VIOS agent, follow these steps:

1. The powervp.x.x.x.x.bff file is in the /Program Files/IBM/PowerVP/PowerVP_Installation/PowerVP_Agent_Installation_Instructions/AIX directory of your workstation, where x.x.x.x is 1.1.3.0 for the current version.

The installp files for IBM GSKit, GSKit8.gskcrypt64.ppc.rte and GSKit8.gskssl64.ppc.rte are also in this directory.

Transfer these files to a directory in your AIX or VIOS system by using File Transfer Protocol (FTP).

To install the AIX and VIOS agent, you must log in to AIX or VIOS shell as root.

2. Run the **installp** command for IBM Global Security Kit (GSKit). If you are installing PowerVP on VIOS, you use the **oem_setup_env** command to enter the unrestricted shell before you run the **installp** command. To support Secure Sockets Layer (SSL), PowerVP includes IBM GSKit and it must also be installed on the partition, even if you do not want to use SSL. To install IBM GSKit, run the following command:

```
installp -acgqwd . GSKit*
```

The installation summary is displayed when this task is complete and you can view a message that indicates the result as SUCCESS for both of the GSKit files. If IBM GSKit was already installed on your partition, a message is displayed.

3. To install PowerVP, run the following command:

```
installp -agXd . powervp.rte
```

The installation summary is displayed when this task is complete and you can view the result as SUCCESS.

4. To complete the configuration of the PowerVP agent, use the iconfig program, which automatically sets the PowerVP configuration file on the server with the values specified. Copy and paste the command as displayed in the browser window and run it on your AIX or VIOS partition. If you are installing the system level agent and are using default ports, you need not use the iconfig program. If you are installing a partition level agent, enter a command similar to this one:

```
/opt/ibm/powervp/iconfig SystemLevelAgent=mysystem.com
```

Where mysystem.com is the host name of your system level agent.

5. If you changed the ports, you must use the iconfig program to set the port for the agent. Enter the following command:

```
/opt/ibm/powervp/iconfig Listen="* 13000"
```

Replace 13000 with the port you selected.

If you do not want to use SSL, the AIX or VIOS agent is configured and ready for use.

If you want to install the AIX or VIOS agent as a partition level agent on other AIX or VIOS partitions, use FTP to transfer the powervp.1.1.3.0.bff file to that system and complete the instructions to install PowerVP on the AIX or VIOS partition.

If you want to use SSL, you must configure SSL on PowerVP by specifying a digital certificate for the AIX or VIOS agent. For instructions, see "Configuring SSL on PowerVP" on page 9.

To start the AIX and VIOS agent, run the following command on your AIX or VIOS partition:

```
/etc/rc.d/rc2.d/SPowerVP
```

This script starts the PowerVP agent as a background process. The log file for the AIX and VIOS agent is in the /var/log/powervp.log file.

Installing the PowerVP GUI application

If you selected to run the PowerVP GUI application in your own application server, you need to deploy the PowerVP GUI application to your application server. The installer copies the PowerVP application to the PowerVP_GUI_Installation subdirectory within the destination folder you selected during the installation. To deploy the PowerVP GUI application to your own application server, you need to consult the documentation for your application server and follow the instructions that it provides to deploy applications.

Installing the IBM i agent

The IBM i agent is automatically installed during installation of PowerVP. The configuration file was set with the values you specified during the installation of PowerVP on your workstation.

If you want to use SSL, you must use Digital Certificate Manager to associate a server certificate with the PowerVP agent. For more information, see "Configuring SSL on PowerVP" on page 9.

If you want to install the agent on other IBM i partitions, you must rerun the installation and select the option to install PowerVP Server Agents only.

To manually install the IBM i agent, use FTP to transfer the SAVSLE00MM.SAVF and the correct SAVSLE00xx.SAVF language file to your IBM i partition to a SAVF object. Then run the Restore Licensed Program (RSTLICPGM) command to install the IBM i agent on the IBM i partition. For example:

- RSTLICPGM LICPGM(5765SLE) DEV(*SAVF) RSTOBJ(*PGM) SAVF(lib/SAVSLE00MM)
- RSTLICPGM LICPGM(5765SLE) DEV(*SAVF) RSTOBJ(*LNG) SAVF(lib/SAVSLE00xx)

Note: xx represents the language code of the system that corresponds to the save file to use.

Installing the Linux agent

Follow these instructions to install the Linux agent:

1. Locate the PowerVP RPM files that are in the /Program Files/IBM/PowerVP/PowerVP_Installation/PowerVP_Agent_Installation_Instructions/LINUX directory
2. Install the following prerequisite utilities on the Linux partition before you install PowerVP:
 - **sysstat**

- **procps**
- **net-tools**
- **ethtool**
- **perf**
- **coreutils**
- **ksh**

The required RPM files depend on the distribution and version of Linux you are using.

- Big endian Red Hat Enterprise Linux (RHEL) or SUSE Linux Enterprise Server (SLES) distributions and versions require the `powervp-1.1.3.1.ppc64.rpm`, `gskcrypt64-8.0.50.42.linux.ppc.rpm`, and `gkssl64-8.0.50.42.linux.ppc.rpm` files.
- Little endian Red Hat Enterprise Linux (RHEL) or SUSE Linux Enterprise Server (SLES) distributions and versions require the `powervp-1.1.3.1.ppc64le.rpm`, `gskcrypt64-8.0.50.42.linux.ppcle.rpm`, and `gkssl64-8.0.50.42.linux.ppcle.rpm` files.
- Ubuntu Linux distributions require the `powervp-1.1.3-1.ppc64le.rpm`, `gskcrypt64_8.0-50.42_ppc64el.deb`, and `gkssl64_8.0-50.42_ppc64el.deb` files.

PowerVP includes a kernel extension module in an RPM file. The Linux distribution name and version are included in the name of the RPM file, so you need to select the correct RPM file for your Linux distribution and install that RPM file. If you do not see your Linux distribution version, you can install the RPM file that contains the `powervp-driver-source` in the name of the file and build the PowerVP kernel extension yourself. For more information, see step 7 on page 9.

3. Use FTP to transfer the necessary RPM files to a directory in your Linux partition, and then run the following commands:
 - For Ubuntu Linux distributions, run the following commands:
 - a. **`dpkg -i gskcrypt64_8.0-50.42_ppc64el.deb gkssl64_9.0-50.42_ppc64el.deb`**
 - b. **`alien -ic powervp-driver-1.1.3-1.ubuntu15.04-3.19.0-15.ppc64le.rpm`**
 - c. **`alien -ic powervp-1.1.3-3.ppc64le.rpm`**

For big endian and little endian Red Hat Enterprise Linux (RHEL) or SUSE Linux Enterprise Server (SLES) distributions and versions, run the following command:

- **`rpm -i powervp-1.1.3.1-1.ppc64.rpm gskcrypt64-8.0.50.42.linux.ppc.rpm gkssl64-8.0.50.42.linux.ppc.rpm powervp-driver-xxx.rpm`**

Replace `powervp-driver-xxx.rpm` with the correct file name for the Linux distribution that is installed on your system.

4. Use the `iconfig` program to set the values in the configuration file. You can run the `iconfig` program to automatically set the PowerVP configuration file on the partition with the values specified. Copy and paste the command as displayed in your browser and run it on your Linux partition. If you are installing a partition level agent, run a command similar to this command:

```
/opt/ibm/powervp/iconfig SystemLevelAgent=mysystem.com
```

Where `mysystem.com` is the TCP/IP host name of your system-level agent.

5. If you changed the ports, use the `iconfig` program to set the port for the agent. Enter the following command:

```
/opt/ibm/powervp/iconfig Listen="* 13000"
```

Replace 13000 with the port you selected.

6. If you do not want to use SSL, the Linux agent is now configured and you can start it by entering the following command:

```
service ibm_powervp start
```

That command starts the PowerVP agent as a service.

If you want to use SSL, configure the SSL on PowerVP by specifying a digital certificate for the agent. For instructions, see “Configuring SSL on PowerVP” on page 9.

7. If you have a different installation or version of Linux, use the `powervp-driver-source-1.1.3.1-1.ppc64.rpm` RPM file to install the source code for the kernel extension module that is required by PowerVP.

For information about how to install the source code for the kernel extension module, see “Additional instructions for Linux.”

Additional instructions for Linux

If you have a different installation or version of the Linux operating system, you need to use the `powervp-driver-source-1.1.3.1-1.ppc64.rpm` RPM file, which installs the source code for the kernel extension module that is needed by PowerVP.

First, you need to install the source code RPM by entering the following command:

```
rpm -i powervp-driver-source-1.1.3.1-1.ppc64.rpm
```

The files are installed in `/opt/ibm/powervp/driver-source`. From that directory, run the **make** command to build a `powervp-driver` RPM file for your Linux system. Note this command requires packages that are needed when building kernel modules.

After you rebuild your kernel to include the PowerVP kernel module, you can then finish the PowerVP installation by following the installation instructions.

Configuring SSL on PowerVP

To use Secure Sockets Layer (SSL) with PowerVP, you must configure SSL on the agent partition.

To use SSL with the PowerVP server, you must have a server certificate. The following instructions assume your server certificate is in PKCS#12 format in a file called `myserver.p12`. You must also know the password used to protect this file and the label for the certificate. The label may have been set when the certificate was created or exported to the PKCS#12 file. You can view the PKCS#12 file to determine the label. You must also determine which port the PowerVP server uses to accept SSL connections and determine whether you want to accept non-SSL connections on a separate port.

Configuring SSL on AIX, VIOS, or Linux

Perform the following steps to configure SSL on your AIX, VIOS, or Linux partition.

1. Copy your server certificate file to the `/etc/opt/ibm/powervp/certs` directory.

2. To determine the certificate label, run the following command:

```
/usr/opt/ibm/gsk8_64/bin/gsk8capicmd_64 -cert -list all -db myserver.p12
```

For Linux, the `gsk8capicmd_64` command path is: `/usr/local/ibm/gsk8_64/bin/gsk8capicmd_64`.

Enter the file password and then you can view the following:

Certificates found

```
* default, - personal, ! trusted, # secret key
```

```
! LOCAL_CERTIFICATE_AUTHORITY_0288C5554(1)
```

```
- MYSYSTEM
```

Your server certificate will be flagged as either `'*` (default) or `'-` (personal). In this example, `MYSYSTEM` is the label of the server certificate.

3. You can set the password of the certificate file (`myserver.p12`) in the PowerVP configuration file (`powervp.conf`) or it can be encrypted in a separate stash file. To stash the password, run the following command:

```
/usr/opt/ibm/gsk8_64/bin/gsk8capicmd_64 -keydb -stashpw -db myserver.p12
```

For Linux, use `/usr/local/ibm/gsk8_64/bin/gsk8capicmd_64`.

Enter the file password and a stash file containing the encrypted password will be created in the same directory as the myserver.p12 file with a .sth extension (for example: myserver.sth).

4. Update the PowerVP server configuration with the certificate file information. For example:
/opt/ibm/powervp/iconfig KeyringFile=myserver.p12 StashFile=myserver.sth CertificateLabel=MYSYSTEM or if you did not create a stash file, enter the following command:
/opt/ibm/powervp/iconfig KeyringFile=myserver.p12 KeyringPassword=yourpassword CertificateLabel=MYSYSTEM

If you store the password in the PowerVP configuration file, it's recommended that you also remove public read access from the /etc/opt/ibm/powervp/powervp.conf file. If your certificate label contains spaces, enclose the label in quotes in the previous commands (for example, **CertificateLabel="My Server"**).

The default PowerVP server ports are 13000 for non-SSL connections and 13001 for SSL connections. To update the PowerVP server configuration for both non-SSL and SSL connections on the default port enter the following command:

```
/opt/ibm/powervp/iconfig Listen="* 13000", "* 13001 ssl"
```

To update the PowerVP server configuration for only SSL connections on the default port, enter the following command:

```
/opt/ibm/powervp/iconfig Listen="* 13001 ssl"
```

Note: You can also edit the configuration file at /etc/opt/ibm/powervp/powervp.conf. The configuration file can also be edited to control the SSL protocols (TLS 1.0, TLS 1.1, and TLS 1.2) and Transport Layer Security (TLS) cipher suites that will be supported by the PowerVP server.

5. Start (or restart) the PowerVP server.

Configuring SSL on IBM i

Follow these steps to complete the SSL configuration on your IBM i partition.

1. Copy your server certificate file to the **/etc/opt/ibm/powervp/certs** directory.
2. By using the IBM i Digital Certificate Manager (DCM) web interface, import your server certificate into the *SYSTEM certificate store.
3. By using DCM, update the certificate assignment for the IBM PowerVP Server (application ID QIBM_QPF_POWERVP_SERVER) and select your server certificate.

Note: By using DCM, you can also update the application definition to select the SSL protocols (TLS 1.0, TLS 1.1, and TLS 1.2) and TLS cipher suites that will be supported by the PowerVP server.

4. The default PowerVP server ports are 13000 for non-SSL connections and 13001 for SSL connections. To update the PowerVP server configuration for both non-SSL and SSL connections on the default ports, enter the following command:

```
CALL QSLE/QPFICONFIG PARM('Listen="* 13000", "* 13001 ssl"')
```

To update the PowerVP server configuration for only SSL connections on the default port, enter the following command:

```
CALL QSLE/QPFICONFIG PARM('Listen="* 13001 ssl"')
```

Note: You can also edit the configuration file at /QIBM/UserData/PowerVP/powervp.conf.

5. Start (or restart) the PowerVP server by using the following command:
STRTCPSVR SERVER(*POWERVP)

Configuring SSL on partition-level agents

All of the partition level agents on a system register with their system-level agent by connecting to the system-level agent. The partition level agent must be able to connect to one of the listening addresses and

ports of the system level agent. If the partition level agent must use SSL to connect to the system level agent (for example, the system level agent allows only SSL connections), the partition level agent must be configured for SSL and its keyring file must include the certificate authority (CA) that has issued the system level agent's certificate. If both systems' certificates were issued by the same CA, the keyring file will already contain the correct CA.

Starting and using the PowerVP GUI

PowerVP is a real-time monitor that can collect and update the performance information as frequently as every second. The PowerVP GUI displays the real-time data.

Configuring the PowerVP agent

The PowerVP configuration file `powervp.conf` contains the directives that tell the PowerVP agent how it is to run. If your partition has multiple host names or IP addresses, you might need to tell PowerVP which host name and IP address it is to use. You can specify a host name and IP address by using the **AgentHostNameAndIP** directive for PowerVP to use for the agent. If you do not provide this directive, PowerVP uses the first host name that the system returns on the host name `api` for the agents host name and it also uses the first IP address for this host name. If you want the agent to use another host name or IP address, specify that using the **AgentHostNameAndIP** directive as follows:

```
AgentHostNameAndIP name ip
```

Where *name* is the host name and *ip* is the IP address.

Starting the PowerVP GUI

You can start the PowerVP GUI from the **Start** menu in the Windows operating system or one of the other methods for starting the PowerVP GUI, such as on the desktop, in the Quick Launch bar, or in a Program Group. Locate the PowerVP icon to start the PowerVP GUI.

If you deployed the PowerVP GUI application to an application server somewhere in your network, you need to navigate your browser to a URL on the application server. This process is often done through a URL link on a web page. Contact the person that installed PowerVP for the URL to use to run PowerVP.

Connecting to the agent

Connect PowerVP with the PowerVP agent on the POWER system. After you connect PowerVP to the PowerVP agent, PowerVP can display POWER system information in the system information and host information sections.

Perform the following steps to connect PowerVP to the PowerVP agent:

1. From the PowerVP welcome page, click **New Connection**. The New Connection page is displayed.
2. Complete the information in the New Connection page.
 - In the **Host name** field, enter the TCP/IP host name of the POWER system where you installed the system level agent. You must connect to the system level agent and enter a valid user name and password on the POWER system.
 - If you did not use the default ports, you must change the ports.
 - If you chose to use only SSL during installation of the PowerVP agent, select the **Secure** check box and clear the **Non-secure** check box.
 - If you want to view current data, clear the **Load Date** check box.
 - If you want to view historic data from the PowerVP data store, select the **Load Date** check box and specify a date and time.
3. Click **Connect**.

PowerVP is now connected to your POWER system and your POWER system information is displayed in the system information and host information sections. The main panel displays the nodes and partitions on your POWER system. If you are playing historic data from the data store, you can use the DVR buttons to fast forward, rewind, pause, and so on.

Thresholds and alerts

You can configure the PowerVP agent to monitor the CPU and bus utilization levels and to generate an alert if your IBM POWER8 processor-based system's utilization exceeds a configured threshold for a configured time period.

Before using this feature, analyze your system by using PowerVP graphical user interface (GUI) to determine the system's utilization characteristics. You can use this information to set up thresholds for alerts.

An *alert* is a message to syslog on AIX Virtual I/O Server (VIOS) and Linux operating systems and a message to the system operator's message queue, QSYSOPR on IBM i. You can then use the operating system mechanisms to watch for the message and inform your system administrators about the issue. The configuration parameters consist of a threshold (utilization percent) that triggers the alert, an amount of time (duration) for the utilization to exceed that threshold before an alert is generated (so any minor performance spikes can be avoided), an amount of time to wait before it sends another alert (realert) while the threshold is still being exceeded, and a syslog level (for AIX, VIOS, and Linux). The realert time allows your system administrators to take action before another message is generated.

You can then use applications that you might already have that monitor QSYSOPR or syslog to generate text messages, e-mails, or pages, to inform system administrators about the situation.

The system CPU utilization and the individual partition CPU utilization can be monitored. You can also monitor the power bus utilizations, internode (A) bus, intranode (X) bus, memory controller bus, and the I/O bus.

The `powervp.conf` PowerVP agent configuration file contains the configuration information. This file is located in the `/etc/opt/ibm/powervp` directory on AIX, VIOS, and Linux and in the `/QIBM/UserData/powervp` directory on IBM i. The configuration file directives follow:

- **UtilizationAlertPartitionCPU** *percent duration realert level*
- **UtilizationAlertSystemCPU** *percent duration realert level*
- **UtilizationAlertAbus** *percent duration realert level*
- **UtilizationAlertXbus** *percent duration realert level*
- **UtilizationAlertMCbus** *percent duration realert level*
- **UtilizationAlertInputIObus** *percent duration realert level*
- **UtilizationAlertOutputIObus** *percent duration realert level*

In each of these directives, the *percent* refers to the utilization percent threshold at which alert monitoring is started. The *duration* is the amount of time, in seconds, for which the utilization must exceed the specified *percent* before the alert is generated. The *realert* is the amount of time, in seconds, after an alert has been sent before another alert will be sent (provided the utilization continues to exceed the threshold). The *level* specifies the severity of the error to be reported. The *level* is used only on AIX, VIOS, and Linux. It is ignored on IBM i. The *level* defaults to *Notice* for Linux, AIX, and VIOS. Valid values for *level* include the syslog severity levels of *Emergency*, *Alert*, *Critical*, *Error*, *Warning*, *Notice*, and *Informational*. The syslog Facility will be *daemon*.

You can choose to configure only the situations you want to monitor. For situations that you do not want to monitor, do not provide any configuration file directive.

The SystemCPU utilization is the CPU utilization of the entire POWER processor-based system which includes all partitions for all activated cores.

The PartitionCPU utilization is the CPU utilization of an individual partition based on its entitled processor capacity.

The Abus and Xbus utilization can be used to monitor the partition affinity of the system.

The MCBus utilization can be used to monitor memory allocation and affinity of the system.

The InputIObus and OutputIObus utilization allow you to separately monitor the incoming and outgoing I/O for high utilization.

For the IBM i operating system, PowerVP sends a message to the QSYSOPR message queue. The message IDs follow:

- SLE0121 for UtilizationAlertSystemCPU
- SLE0122 for UtilizationAlertPartitionCPU
- SLE0123 for UtilizationAlertAbus
- SLE0124 for UtilizationAlertXbus
- SLE0125 for UtilizationAlertMCbus
- SLE0126 for UtilizationAlertInputIObus and UtilizationAlertOutputIObus

For the AIX, VIOS, and Linux operating system (and for IBM i) the message text starts with a message identifier:

- MSG0107 for UtilizationAlertSystemCPU
- MSG0106 for UtilizationAlertPartitionCPU
- MSG0108 for UtilizationAlertAbus
- MSG0109 for UtilizationAlertXbus
- MSG0110 for UtilizationAlertMCbus
- MSG0111 for UtilizationAlertInputIObus and UtilizationAlertOutputIObus

In addition to a message in QSYSOPR or syslog, PowerVP logs the message in the job log for the PowerVP agent on IBM i and in the `/var/log/powervp.log` file on AIX, VIOS, and Linux.

Permanent data store

The PowerVP agent can be configured to write the PowerVP performance data to a file, also known as the *data store*. You can configure the data store to enable logging of PowerVP performance data, the path name of the file where the data is stored, how long the current file is used before it is closed and a new data store file is started or how large the file is before it is closed and a new data store file is started. You can also configure how long the archived files are to be retained on your POWER processor-based server.

The PowerVP agent can be configured to write the PowerVP performance data to a file called the *data store*. To configure the data store, you must change the configuration file for your PowerVP agent. When you are recording PowerVP performance data, PowerVP uses space on your POWER system disk. The amount of space depends on the size of your POWER system, the number of partitions that are configured on your POWER system, and the sample interval that is configured for the PowerVP agent. You must monitor the size of the PowerVP recording file to ensure sufficient space exists for the file. You can use the `LogFileRotation` and `LogFileArchive` directives to control the PowerVP data store. You can also use the `SampleInterval` directive to decrease the amount of performance data recorded in the data store. The PowerVP configuration file is `/etc/opt/ibm/powervp/powervp.conf` on AIX, VIOS, and Linux operating systems and `/QIBM/UserData/powervp/powervp.conf` on the IBM i operating system.

Configuration file directives

You can enable or disable logging, specify a path name for the data store file, specify the file size or time frame for rotating to a new file, and set the number of days for archive retention by using the configuration file directives that are associated with the PowerVP data store.

LogData

The LogData configuration file directive is used to enable or disable the PowerVP data store. This directive has one parameter with values *Yes* or *No*. The default value is *No*.

LogFilePath

The LogFilePath configuration file directive is used to specify where PowerVP must store the data store file. The default location is `/opt/ibm/powervp/logs` on AIX, VIOS, and Linux and `/QIBM/UserData/powervp/logs` on IBM i. Do not store any other files in this folder. The PowerVP agent can be configured to archive old data store files. The archival efficiency is better if only PowerVP data store files are stored in the folder. The file name is controlled by PowerVP and it has the following format: **PVPmmdyyyyhhmmss.csv**, where *mmdyyyy* is the month/day/year and *hhmmss* is the hour/minute/second of the first performance data in the file. The file is in the comma separated value (CSV) format, which you can import into spreadsheet applications.

LogFileRotation

The LogFileRotation configuration file directive allows you to control when PowerVP closes the current data store file and starts storing data in a new file, allowing you to control the size of the data store files. You can set the number of hours that PowerVP writes to a file before it closes the file and starts writing to a new file. You can also set the size of the data store file. To specify a time, provide a number between 1 and 24 followed by the letter H. For example, 12H means 12 hours. To specify a file size, provide a number followed by either the letter M (for megabytes) or G (for gigabytes). If you use a file size, PowerVP performs the file rotation when the file size is close to this value. Because each line of the file is a different size, the resulting file sizes will not match this value. The minimum value for size is 100M. PowerVP rotates the log file at midnight, 00:00:00 time for either rotation method.

LogFileArchive

The LogFileArchive configuration directive allows you to control how long the data store files are retained on your POWER server. This directive takes a single numeric value, which is the number of days to retain the data store files. At midnight, the PowerVP agent determines if any files must be archived and performs the archive operation. The default value is 7 days. If the data store file is being used by another application when PowerVP attempts to archive it, the file is marked for deletion and will be deleted when that application closes the file.

To play the recorded information from the PowerVP data store, see “Starting and using the PowerVP GUI” on page 11. You provide a date and time as the start time, which instructs the PowerVP agent to locate the data for that date and time and sends it to the PowerVP GUI.

Data store format

The data store is written with unique records for each type of performance data with comma-separated values (csv). For the permanent data store to be used by other tools, it is structured as information entries for each different record that contains performance data. These information entries are written at the start of every file. These records provide column headings for the row type. The first column of each row is the *key* for the data in that row. All rows that match the first column (except for the AAA rows) are the same data, but refers to a different time or different hardware processor module, hardware processor core, partition, disk, Ethernet line, or event.

The special row types are: AAA, and ZZZZ. The AAA rows contain general information about the POWER system and about PowerVP. These row types include the PowerVP version, operating system type and release level for the partition the agent is running on, the sample rate, and other information about the POWER system. The ZZZZ rows contain the time stamp information. The time stamp information has a time stamp identifier in the second column that is in the format Txxxxx where x is

numeric, and T00000 is used for the information row. The ZZZZ rows also have a column for the military time, date, numeric time stamp (seconds since 00:00 hours, Jan. 1, 1970, UTC) and the numeric timebase register from the POWER system. The time stamp identifier is used in all of the other entries to refer to the time for the data in the row.

The second column of all rows (except the time stamp rows and the information rows) is the occurrence count column. An occurrence count of zero means that the rest of the row is column labels, while the rest of the rows are data snapshots with the occurrence count incremented for unique information for that time. The third column of every row (except the time stamp rows and the information rows) is the time column and is used to correlate the row to the correct time stamp row.

Timestamp entries

Timestamp definition record also has a definition row with the second column set to T00000, which indicates that it is the header row.

- ZZZZ, T00000, military time in HH:MM:SS, date in DD-MM-YYYY, timestamp in seconds elapsed since 00:00 hours on Jan 1 1970 UTC, timebase register

Timestamp rows are similar to the following examples. Since PowerVP rolls over the data store file at midnight every day, the largest time identifier is T86400.

- ZZZZ, T00001, 13:50:00, 25/01/2014, 1390679400
- ZZZZ, T00002, 13:50:01, 25/01/2014, 1390679401

PowerVP Information entries

The following entries are included as the first entries in every file. The entries are listed only once in the file.

The PowerVP program row contains PowerVP.

- AAA, progname, PowerVP

The PowerVP version row contains the version, release, and modification level (VRM) of the PowerVP product. PowerVP V1R1M3 supports data store. The version is updated with the current VRM and fix package used to produce the data. For example:

- AAA, Version, V1R1M3
- AAA, Version, V1R1M3 SP1

The operating system row contains the operating system for the agent. The file is in ASCII format for all operating systems, including IBM i. The operating system row has one of the following entries:

- AAA, Operating system, AIX
- AAA, Operating system, IBM i
- AAA, Operating system, Linux
- AAA, Operating system, VIOS

The operating system version row contains the version of the operating system the agent was running on. For IBM i, the entry is similar to the following example:

- AAA, Operating system version, V7R2M0

For AIX, the entry is similar to the following example:

- AAA, Operating system version, 7.1.0.0

For Linux, the entry is similar to the following examples:

- AAA, Operating system version, RHEL 7.0
- AAA, Operating system version, SLES 11 SP3

- AAA, Operating system version, VIOS 2.2.3.4

The agent type row indicates whether the row is written by a system level agent or a partition level agent. System level agents write both system and partition data to the file, and partition level agents write only partition data to the file.

- AAA, Agent type, system
- AAA, Agent type, partition

The TCP/IP host name row contains the host name of the TCP/IP agent. The host name might not be fully qualified, depending on the operating system TCP/IP configuration on the agent.

- AAA, Agent host name, mysystem
- AAA, Agent host name, mycompany.com

The processor version row contains the processor version of the POWER system. The entry is similar to the following examples:

- AAA, Processor version, POWER7
- AAA, Processor version, POWER8

The agent sample rate row contains the sample rate in seconds.

- AAA, Sample rate in seconds, 1

The power system clock frequency row contains the clock frequency in megahertz.

- AAA, Clock Frequency in MHz, 3920

The power system timebase frequency row contains the timebase frequency in megahertz

- AAA, Timebase Frequency in MHz, 512

The power system serial number row contains the serial number of the Power system.

- AAA, System Serial Number, 10CD93T

The partition time zone information row contains the partition time zone offset from Greenwich Mean Time (GMT).

- AAA, Partition Timezone offset, 6
- AAA, Partition Timezone offset, -2

Topology

The topology entries contain the physical topology information for the POWER system. Multiple topology entries might exist in the file. If nodes, chips, or cores are activated or deactivated, subsequent topology records are stored in the file that contain the time when the topology change occurred.

System topology:

The system topology SYSTOP record contains integers for the number of nodes, chips, cores, and virtual processors on the POWER system. Each file contains information for only one POWER system, thereby the SYSTOP record contains only 0 and 1 rows. If the hardware configuration changes while PowerVP is running, a new SYSTOP record is written to indicate the new configuration.

- SYSTOP, 0, T00000, number of nodes, number of chips, number of cores, number of virtual processors

After a hardware configuration change, the entry is similar to the following examples:

- SYSTOP, 1, T00001, 3, 12, 96, 384
- SYSTOP, 1, T00090, 4, 16, 128, 512

Chip topology:

The chip topology CHIPTOP rows contain the information about the processor modules (chips) on the system. Each processor module has its own row with the second column that starts with 1 and increments for each processor module. If the topology changes, the CHIPTOP rows are written with the time column set to indicate the time of the change. The rows contain integers for the ID columns and width columns. For the bus columns, the link attached column contains either the Yes or No value. If the bus is connected, the value is Yes. The corresponding link endpoint node or chip ID column indicates the node or chip ID of the endpoint. If the bus is not connected, the value is No. The endpoint ID is 0, which does not indicate that it is connected to chip or node 0. For the GX, PHB, and MC buses, the attached column also has the Yes or No value.

- CHIPTOP, 0, T00000, physical chip ID, hardware chip ID, fabric chip ID, hardware node ID, A (ABC) bus width, X (WXYZ) bus width, GX bus width, MC bus width, PHB bus width, A0 (A) link attached, A0 (A) link endpoint node ID, A1 (B) link attached, A1 (B) link endpoint node ID, A2 (C) link attached, A2 (C) link endpoint node ID, X0 (W) link attached, X0 (W) link endpoint chip ID, X1 (X) link attached, X1 (X) link endpoint chip ID, X2 (Y) link attached, X2 (Y) link endpoint chip ID, X3 (Z) link attached, X3 (Z) link endpoint chip ID, GX0/PHB0 bus attached, GX1/PHB1 bus attached, PHB2 bus attached, PHB3 bus attached, MC0 bus attached, MC1 bus attached, MC2 bus attached, MC3 bus attached
- CHIPTOP, 1, T00001, 0, 0, 0, 0, 2, 2, 2, 2, 2, Yes, 2, Yes, 2, No, 0, Yes, 1, Yes, 2, No, 0, No, 0, Yes, No, No, No, No, Yes, No, No
- CHIPTOP, 002, T00001, 1, 1, 1, 0, 2, 2, 2, 2, 2, Yes, 2, Yes, 0, No, 0, Yes, 1, Yes, 2, No, 0, No, 0, No, Yes, No, No, Yes, No, No, No

Core topology:

The core topology CORETOP records contain the information about the processor cores on the system. A second column of 0 indicates a column heading record. Each core has its own record with the second column that starts with 001 and increments for each core. If the topology changes, the CORETOP records will be rewritten after the new CHIPTOP records are written with the time column set to indicate the time of the change. The records contain integers for all items except the core state, which is a character value, and it has one of the following values: Not installed, Guarded off, Unlicensed, Shared, Borrowed, or Dedicated. The assigned partition is integer 65535 if the core is not assigned to a dedicated partition or if it is in the shared group.

- CORETOP, 0, T00000, core ID, chip ID, module ID, node ID, core state, logical processor ID, processor identification register, assigned partition ID, nominal frequency in MHz, current frequency in MHz, primary affinity domain, secondary affinity domain
- CORETOP, 1, T00001, 0, 0, 0, 0, Dedicated, 1, 0, 1, 4000, 4004, 0, 1
- CORETOP, 2, T00001, 1, 0, 0, 0, Shared, 4, 4, 65535, 4000, 4004, 0, 1
- CORETOP, 3, T00001, 2, 0, 0, 0, Borrowed, 8, 8, 65535, 4000, 4004, 0, 1

Registered partitions:

The registered partitions REGLPARS entries define the partitions that are registered as partition level agents with the system level agent. A second column of 0 indicates a column heading record. Each registered partition entry has its own record with the second column starting with 001 and increments for each registered partition. If a partition registers or removes itself from the partition list, the REGLPARS records are rewritten with the time column set to indicate the time of the change.

- REGLPARS, 0, T00000, partition ID, version, operating system, agent type, processor version, authentication type, host name, IP address
- REGLPARS, 1, T00001, 0, 3, AIX, partition, POWER8, system, myvios.domain.com, 9.5.11.11
- REGLPARS, 2, T00001, 1, 3, IBM i, partition, POWER8, system, myibmi.domain.com, 9.5.11.12
- REGLPARS, 3, T00001, 2, 3, Linux, partition, POWER8, system, myrhel.domain.com, 9.5.11.13

- REGLPARS, 4, T00001, 3, 3, VIOS, system, POWER8, system, myvios.domain.com, 9.5.11.14

Affinity information by domain:

The affinity domain AFFDTOP topology entries define the affinity domains on the system. A second column of 0 indicates a column heading record. Each affinity domain has its own record with the second column that starts with 001 and increments for each affinity domain. If the topology changes, the AFFDTOP records are rewritten with the time column set to indicate the time of the change.

- AFFDTOP, 0, T00000, primary domain, secondary domain, total processor units, free dedicated processor units, free shared processor units, total memory, free memory, number of partitions in domain
- AFFDTOP, 1, T00001, 1, 0, 600, 100, 100, 256, 10, 13

Affinity information by partition:

The affinity partition AFFPTOP topology entries define the affinity by partition. A second column of 000 indicates a column heading record. Each partition has its own record with the second column that starts with 001 and increments for each partition. If the topology changes, the AFFPTOP records are rewritten with the time column set to indicate the time of the change. The placement spread field is a character and contains one of the following values: unknown, contained in primary domain, contained in secondary domain, spread across secondary domain, wherever fits, or scramble. The other fields are all numeric. Each partition record also contains a number of affinity element records. The last field in the AFFPTOP record indicates how many affinity element records follow and the partition ID field in the elements matches the partition ID in the partition record.

- AFFPTOP, 0, T00000, partition ID, assignment order, placement spread, affinity score (0-100), number of affinity domain elements
- AFFPTOP, 1, T00001, 1, 1027, 1, 90, 2

The affinity elements are similar to the following examples:

- AFFPELE, 0, T00000, partition ID, primary affinity domain, secondary domain index, dedicated processor units allocated, dedicated memory allocated default, dedicated memory allocated reserved 1, dedicated memory allocated reserved 2, dedicated memory allocated 16 GB pages
- AFFPELE, 1, T00001, 1, 2, 2, 2, 600, 124, 0, 0
- AFFPELE, 1, T00001, 1, 2, 3, 2, 600, 244, 0, 0

Affinity domain information by virtual processor:

The AFFVPROC affinity domain information that is provided by the virtual processor records define the domain affinity by virtual processor. A second column of 000 indicates a column heading record. Each virtual processor has its own record with the second column that starts with 001 and increments for each virtual processor. If the topology changes, the AFFVPROC records are rewritten with the time column set to indicate the time of the change. All fields are numeric.

- AFFVPROC, 0, T00000, partition ID, virtual processor index, physical processor index, primary affinity domain index, secondary affinity domain index
- AFFVPROC, 1, T00001, 1, 3, 4, 2, 4

System Agent Data

The system level agent information is retrieved from the hypervisor for all partitions and cores on the system.

CPU Utilization by core:

The CPU utilization data is retrieved from the POWER system hypervisor by core ID. Each core has its own row with the second column that starts with 1 and increments for each core.

- SCPUBC, 0, T00000, core ID, user plus kernel Processor Utilization Resource Register (PURR) delta, unfiltered PURR delta, run instructions delta, total run cycles delta, timebase delta for this sample, current core frequency in MHz
- SCPUBC, 1, T00001, 8, 105582, 253924, 15066481, 98522555, 180168156, 3255

Timebase cycles by core:

The time base cycle utilization is retrieved by core. Each core has its own row with the second column that starts with 1 and increments for each core.

- SCYCBC, 0, T00000, core ID, time base cycles delta, timebase delta
- SCYCBC, 1, T00001, 8, 239116358, 527434092

Bus utilization by Timebase cycles by core:

The POWER bus utilization is retrieved by chip. Each chip has its own row with the second column that starts with 1 and increments for each chip. POWER8 systems do not provide this information. Therefore the file does not contain any records for POWER8.

- SBUSBCH, 0, T00000, chip ID, A0 (A) bus utilization percent, A1 (B) bus utilization percent, A2 (C) bus utilization percent, X0 (W) bus utilization percent, X1 (X) bus utilization percent, X2 (Y) bus utilization percent, X3 (Z) bus utilization percent, inbound PHB0 or GX0 bus utilization percent, inbound PHB0 or GX0 bus rate, inbound PHB1 or GX1 bus utilization percent, inbound PHB1 or GX1 bus rate, inbound PHB2 bus utilization percent, inbound PHB2 bus rate, inbound PHB3 bus utilization percent, inbound PHB3 bus rate, outbound PHB0 or GX0 bus utilization percent, outbound PHB0 or GX0 bus rate, outbound PHB1 or GX1 bus utilization percent, outbound PHB1 or GX1 bus rate, outbound PHB2 bus utilization percent, outbound PHB2 bus rate, outbound PHB3 bus utilization percent, outbound PHB3 bus rate, MC0 bus utilization percent, MC1 bus utilization percent, MC2 bus utilization percent, MC3 bus utilization percent
- SBUSBCH, 1, T00001, 5, 1, 0, 0, 4, 3, 0, 1, 65, 0, 0, 0, 0, 0, 0, 5, 389, 0, 0, 0, 0, 0, 9, 8, 0, 0

CPU utilization by partition:

The CPU utilization by partition consists of the entitled, capped, uncapped, donated, and idle CPU cycles by partition. Each partition has its own row with the second column starting with 1 and increments for each partition.

- SCPUBP, 0, T00000, partition ID, entitled capacity, capped cycles, uncapped cycles, donated cycles, idle cycles
- SCPUBP, 1, T00001, 6, 100.00, 51.79, 41.79, 0.0, 2.71

Partition metrics:

The partition metrics record consists of individual partition metrics. Each partition has its own row with the second column starting with 1 and increments for each partition.

- SMETRICBP, 0, T00000, partition ID, performance data version, delta timebase cycles waiting for entitlement, delta number of times waited on entitlement, delta timebase cycles waiting for a physical processor, delta number of times LPAR dispatched to run, delta home processor dispatches, delta primary affinity domain dispatches, delta secondary affinity domain dispatches, delta remote dispatches, delta dedicated donate processor dispatches, instruction count delta, timebase cycles delta
- SMETRICBP, 1, T00001, 6, 8, 300, 2, 200, 20, 18, 1, 1, 0, 0, 23245, 34323895

Partition Agent Data

The following records are partition level agent records and therefore are collected by all agents because the system level agents collect partition level information about themselves.

Partition status:

The partition status row contains partition status information. To retrieve this information on the AIX operating system, PowerVP uses the perfstat() application programming interfaces (APIs); on the Linux operating system, it uses information from the /proc/ppc64/lparcfg file; and on IBM i, it uses system APIs to get the information.

- PSTAT, 0, T00000, partition ID, dedicated or shared, capped or uncapped, donate enabled, entitled capacity, active processors in shared pool, partition name
- PSTAT, 1, T00001, 4, Shared, Uncapped, No, 1, 7, PARTITION4

Partition cycles per instruction (CPI):

The partition CPI row contains the cycles per instruction information obtained from the performance monitoring unit (PMU). PowerVP collects event counts for the following groups on POWER8:

pm_utilization, pm_cpi_stack2, pm_cpi_stack4, pm_cpi_stack15, pm_cpi_stack18, pm_dsource1, pm_dsource4, pm_dsource5, pm_dsource6, pm_dsource7, and pm_dsource8; and for the following groups on POWER7: pm_dlatencies3, pm_cpi_stack1, pm_cpi_stack2, pm_cpi_stack7, pm_dsource1, pm_dsource2, pm_dsource3, pm_dsource4, pm_dsource5, pm_dsource6, pm_psource10, pm_dsource12, and pm_prefetch2.

- PCPI, 0, T00001, group name, event 1 count, event 2 count, event 3 count, event 4 count, event 5 count, event 6 count
- PCPI, 1, T00001, pm_utilization, 9790840, 8165099, 3656739, 2536237, 2536237, 8165035

Partition Ethernet throughput:

The partition Ethernet throughput row contains the Ethernet send/receive counts per Ethernet interface on the partition.

- PENET, 0, T00000, interface name, kilobytes sent, kilobytes received
- PENET, 1, T00001, en0, 6256, 3425

Partition disk I/O:

The partition disk I/O row contains the disk I/O per disk for the partition.

- PDISK, 000, T00000, name, kilobytes read, kilobytes written
- PDISK, 001, T00001, hdisk0, 2.57, 4.75

Partition CPU cycles:

The partition CPU cycles row contains the types of CPU cycles that are used by the partition.

- PCPU, 0, T00000, entitled cycles, capped cycles, uncapped cycles, donated cycles, idle cycles
- PCPU, 1, T00001, 100.00, 51.47, 43.27, 0.0, 2.71

Recording and playing performance metrics in PowerVP

You can use PowerVP to record and save performance metrics. The PowerVP data store on the agent can continually record your system's PowerVP performance metrics. Alternatively, you can record and save the performance metrics to a file on your application server by using PowerVP.

If the PowerVP data store on the agent records your system's PowerVP performance metrics, you can replay the performance metrics from any time period that is contained in the data store. When you replay historic performance data from the server, you can use the DVD buttons to fast forward, rewind, pause, and so on.

You can use PowerVP to record benchmark performance metrics of your system while it is running a specific workload. This benchmark recording can then be used to compare to a system that is running the same workload with changes to the system configuration. If you are planning to change your system configuration, use the recording function to determine the performance impact of your configuration changes. You can compare the recorded system performance metrics with your current system performance by playing the recording in a web browser and viewing the current system performance metrics in another web browser. To view recorded system performance metrics and the current system performance metrics simultaneously, you must open the PowerVP GUI in two separate web browser windows. You cannot view recorded system performance metrics and the current system performance metrics simultaneously in the same PowerVP GUI instance.

You can replay the recorded performance metrics to view saved performance information. You can use the saved performance information to verify the impact of configuration changes to a workload. You might record the current workload, then change your settings or environment, and record the same workload again. You can then compare the two recordings to verify whether your changes to the settings or environment are successful.

For recording over long periods of time, you might want to change the sample interval on your PowerVP agent. The default is 1 second. If you are going to save recordings, it is recommended that you change this interval to a larger value. The smaller the interval, the more data is recorded and the file grows much faster. The configuration file is in `/QIBM/UserData/PowerVP/powervp.conf` on IBM i and in `/etc/opt/ibm/powervp/powervp.conf` on AIX/VIOS and Linux. The directive name is `SampleInterval`. If you change this value, you need to stop and restart the PowerVP agent in order for the new value to be used.

Recording performance metrics

To record performance metrics, click **Start Recording** in the DVR area of the PowerVP GUI. To record the partition-specific performance metrics also, click **Start recording** on the partition page.

To stop recording performance metrics, click **Stop Recording**.

Replaying recorded performance metrics

To view recorded performance metrics, disconnect the PowerVP GUI from any running system and load a recorded performance metrics file. To load a recorded performance metrics file, click **Load** to display a Load panel, which shows a folder for each system you have recorded. Select the file with the name containing the date and time of the recording.

Click **Select** and the recorded performance metrics start to play. You can use the buttons in the **Playback** section to control the playback. Standard playback controls are available: Rewind, Pause, Fast Forward, and Stop. You can open the nodes to view the core information.

If you recorded partition-specific performance data at the same time, that data is contained in a separate file and must be loaded in another web browser that contains the PowerVP GUI. To view partition-specific performance data, open a web browser with an instance of the PowerVP GUI and click **Load**. The files with the `.slt` extension contain system data, and the files with the `.prt` extension contain partition-specific performance data.

When you view recorded performance metrics in the PowerVP GUI, you cannot connect the PowerVP GUI to a running system or partition. To view the performance metrics for a running system or partition, you must open a separate web browser with an instance of the PowerVP GUI and connect PowerVP to the system or partition.

Using PowerVP to access the VIOS Performance Advisor tool

The VIOS Performance Advisor tool provides advisory reports that are based on the key performance metrics on various partition resources that are collected from the VIOS environment.

The VIOS Performance Advisor tool is included with VIOS starting with VIOS version 2.2.2.0. You can configure the PowerVP agent that is running on VIOS to run the VIOS Performance Advisor tool on a schedule by using the following VIOSAdvisor directive:

VIOSAdvisor runTime timeToRun

Where *runTime* is the duration of time that you want VIOS Performance Advisor to collect information before it generates the report. This value must be in the range 10 - 60, in minutes. The *timeToRun* is the hour of the day that you want PowerVP to start the VIOS Performance Advisor. This value must be in the range 0 - 23, in hours, corresponding to the military time of day. If you want to run multiple times per day, you can specify multiple VIOSAdvisor directives.

To display the VIOS Performance Advisor reports in the PowerVP graphical user interface (GUI), you must connect the PowerVP GUI to the VIOS agent. Then, from the **VIOS Performance Advisor** section of the GUI, you can view the most recent report or view a report from a specific date and time. When you choose to view the most recent report, PowerVP finds the report and sends it to the PowerVP GUI. When you specify a date and time, PowerVP sends the report that has the time stamp equal to or the first report greater than the specified date and time. Alternatively, you can create a new report by running the VIOS Performance Advisor tool for a specific duration. When you choose to create a new report, PowerVP starts the VIOS Performance Advisor tool. When the specified duration is complete, PowerVP receives and displays the output from the VIOS Performance Advisor tool.

PowerVP security considerations

The PowerVP GUI prompts for a user ID and password to access the system-level agent partition. Therefore, you must have a valid user ID and password for that partition to use the PowerVP GUI. There are no other specific authorities that are needed to view the PowerVP performance information.

When you navigate into the partition to view the partition-specific data, the PowerVP GUI is connecting to the partition level agent to retrieve that data. Review the following security considerations for the partition level agents.

- PowerVP authenticates you to the system level agent and if your authentication is successful, you can navigate into the partition to view the partition-level data without logging on to the partition-level agent partition.
- Otherwise, you can require that a user ID and password are entered for the partition.

In the `powervp.conf` partition-level agent configuration file, the **ClientAuthentication** directive controls the method of authentication for the partition-level agent. This directive's default value is **system**. This default value means that you must authenticate to the system-level agent partition prior to drilling down into the partition level agent performance information. You can specify **partition** as the value for this directive, which causes the PowerVP GUI to prompt you for a user ID and password for that partition before you can view the partition drill-down data. You can configure the partitions on a system for different authentication methods, allowing extra security for partitions that require it. If you use **system** as the directive value on all of your partitions, you must provide a user ID and password only once.

Since the PowerVP GUI is now an application that is running in an application server, you can use the security mechanisms that are provided by the application server to enforce extra security access to the PowerVP GUI application. You need to refer to the documentation for your application server for information on how to configure this additional security.

Using PowerVP to investigate and optimize performance

PowerVP is an IBM licensed program product that can help you understand and monitor the performance of your IBM Power Systems.

Clients typically understand the performance of a given logical partition with the help of a comprehensive portfolio of OS-based performance tools from AIX, IBM i, and Linux. However, as the Power based systems have evolved, understanding the performance of the entire Power System as it hosts multiple logical partitions has become more complex with the growing popularity of virtualization and cloud computing. PowerVP was created to fill this gap as it monitors and illustrates the performance of an entire system (or frame). PowerVP will allow you to monitor overall performance and allow you to drill down into more detailed hardware and software views to help you identify and resolve performance issues and to optimize the performance of your Power system.

PowerVP illustrates Power Systems hardware topology in conjunction with resource utilization metrics to help you better understand the system. These resources include nodes, processor modules, chips, cores, Power bus links, memory controller links, GX I/O bus, disk drives, Ethernet, etc. These resource utilizations are portrayed using a colorized heat map technique. These colors and thresholds can be customized to suit a specific client's performance requirements in a meaningful way. For example, green may indicate normal, yellow can signal caution, and red may indicate that some resource is extremely busy and may require action. PowerVP also allows mapping between real and virtual processor resources. For example, click on a partition and see which physical cores are associated with that partition.

PowerVP utilizes a drill-down approach for performance analysis. The system-level view will illustrate overall system-level performance with processor modules and inter-node links. Clicking on a specific hardware node will drill into that node showing the utilization of each core and intra-node links. A listing of each partition is present on both views showing the entitlement, utilization, and physical hardware mapping. Clicking on a specific partition will provide its detailed performance statistics. Collectively, these performance metrics can help you optimize performance with resource balancing, improved affinity, and application efficiency.

General terms used with PowerVP

The following terms are used within PowerVP and are described here as they should be interpreted when using PowerVP. They are arranged in a logical reading order.

System

A physical system is the entire Power System, including all resources for CPU, memory, storage, and so on. This physical system can contain one or more partitions, or virtual systems. Some refer to the system as a frame or a CEC. When referring to PowerVP, do not interchange the terms system (physical) and partition (virtual system).

Partition

A logical partition (LPAR) is the division of a system's resources, such that it can run independently with its own operating system. A physical system can have one or more LPARs (virtual systems). These LPARs can be dedicated or shared (capped or uncapped). A hypervisor, like PowerVP, manages these partitions. Refer to the PowerVM topics in IBM Knowledge Center for descriptions of these related terms (virtual system, entitlement, hardware thread, VIOS, dedicated donate, processor folding, partition types, and so on).

Hardware node

Except for the smallest Power Systems, there is a componentization of the physical system into books, drawers, or nodes. For example, Power 770/780 has up to four drawers, Power 795 has up to eight books.

Socket

A socket is a physical connection on a Power System that connects to one processor module. These modules can be either an SCM (single chip module) or a DCM (dual chip module).

Processor Module

A processor module is an orderable physical entity that connects to a socket. These processor modules can be in the format of an SCM or a DCM. With POWER7, these modules contain processor cores, caches, and other components. For POWER7, a DCM implies two processor chips.

Chip A processor chip is physical integrated circuit that contains processor cores, or caches, or both. POWER7 chips contain up to eight cores with on-chip L1, L2, and L3 caches. POWER8 chips contain up to 12 cores with on-chip L1, L2, L3, and L4 caches. This document does not describe all Power System configurations; but there are significant differences between POWER4/5/6/7/8 chips as well as model footprints within those architecture families.

Core A processor core is a single physical processing unit. With POWER7, up to eight of these cores exist on a single chip and with POWER8, up to 12 of these cores exist on a single chip. Each POWER7 core can have up to four hardware threads that are dispatched to it simultaneously using SMT4. While each POWER8 core can have up to eight hardware threads dispatched to it simultaneously using SMT8. These hardware threads can be called logical cores. A system is sometimes referred to by its total number of physical cores, for example, a 64-core system. LPARs can have an entitlement in terms of a number of cores.

CPU CPU is used to collectively refer to the CPU resources (core, socket, chip, system) for an entity (partition, system) when you are referring to its metrics like CPU utilization, CPU time, CPU cycles. The term CPU is not used explicitly as a specific resource name as it is often confusing. Some refer a CPU to a socket, some to a processor module, and some to a processor core.

Utilization

Utilization is a base performance term that is the percentage of time that a resource is busy. It is normally in the form a percentage, typically from 0% to 100%. Of course, some shared LPARs might have a utilization of greater than 100% if it consumes more CPU resources from the shared pool than its entitlement states.

CPU utilization

This term is much more complex than you might expect. It can refer simply to the percentage of time that the CPU resources are busy. However, with the advent of SMT levels (more than one hardware thread that is dispatched to a core), multi-core systems, and complex processor pipes, CPU utilization becomes more complicated. Each operating system might provide and interpret CPU utilization differently. AIX and IBM i provide utilizations that consider SMT levels and hardware thread dispatch conditions. From this, CPU utilization is rendered where a linear relationship is expected between system throughput and CPU utilization. This metric comes with many assumptions (sufficient other resources for that workload to scale, only true for the actual workload that is used to tune the utilization while other workloads might scale differently, and so on). Linux operating systems currently provide CPU utilizations that are based more on occupancy (hardware thread that occupies a given core). The more that you understand about this topic, the more you realize that other metrics are also needed to best understand your system or application (such as scaling characteristics, instructions consumed, run cycles consumed, contention issues).

Powerbus (W, X, Y, Z, A, B, C)

Powerbus links are a set of links or buses within Power Systems. Within PowerVP on a POWER7, those links that are labeled W, X, Y, or Z are links within a hardware node; those links that are labeled A or B are links between hardware nodes. On a POWER8, those links that are labeled X0, X1, X2, and X3 are links within a hardware node; those links that are labeled A0, A1, or A2 are links between hardware nodes. These Powerbus links carry data between a given chip and other

resources outside that chip (cache, memory, I/O). PowerVP portrays these links and their utilizations. Having a higher Powerbus utilization implies that there is a higher rate of data transfer.

Memory Controller (MC)

The memory controller is a set of links that connects the memory to the socket. The MC buses carry data between the memory controller and the processor. PowerVP monitors and displays these links and their utilizations. A higher MC bus utilization implies that there is a higher rate of data transfer.

I/O (GX or PHB) bus

The I/O bus is a set of links or buses within a Power System that connect the I/O subsystems to the chip. On a POWER7, they are labeled GX; and on a POWER8, they are labeled PHB. These links carry data for storage I/O and network I/O. PowerVP portrays these links, their utilizations, as well as their inbound/outbound data rate. Having a higher GX/PHB bus utilization implies that there is a higher rate of data transfer.

Cycles per instruction (CPI)

CPI is a standard measurement of application efficiency. It is the number of cycles consumed divided by the number of (machine) instructions completed. Normally, a lower CPI is better than a higher CPI. A CPI can be measured for a core, a processor module, a hardware node, or an LPAR with PowerVP. From an LPAR perspective, you can break down CPU utilization, into CPI components (for example, load/store unit, fixed point, global completion table).

CPI stack analysis

CPU utilization can be broken down into CPI components. Load/Store Unit (LSU CPI) reflects the cycles consumed for accessing data (L1, cache, L2 cache, L3 cache, memory). Fixed Point (FXU CPI) reflects cycles that are consumed on running fixed point. Global Completion Table (GCT CPI) reflects cycles that are consumed waiting on the global completion table for pipelining out-of-order instruction execution. PowerVP analysis typically focuses on LSU CPI.

LSU CPI stack analysis

Normally the largest component of CPU utilization is the LSU CPI for OLTP applications. In other words, accessing data consumes most of the CPU resources. A characterization of the time that it accesses data from L1 cache, L2 cache, L3 cache, and memory; this also notes whether the accesses are for cache/memory for a chip or for another chip on the same processor module or hardware node or distant hardware node. On a POWER8, a L4 cache is also used.

PowerVP FAQ

Once you have PowerVP (Power Virtualization Performance) installed and monitoring your Power Systems, this investigation Q&A section helps you better understand the information available and optimize the performance of your systems. It attempts to help you define criteria for your utilization thresholds and provide some best practices to help you interpret the results and optimize your system performance.

When should I monitor with PowerVP?

You should proactively use your performance management tooling to understand your system(s) performance. It is best to have baseline information that reflects current performance levels. If you try to optimize performance later, you will have a baseline to compare to your improvement attempts. It is also good to have a baseline for a time period where the system performance is normal to compare to in case you have a performance issue in the future. Ideally you could run PowerVP all the time. You can configure the PowerVP agent to "record" the PowerVP performance information in a file on the Power system to accomplish this. Remember you can monitor only (whether real time or with playback) information that was recorded with PowerVP as it cannot use historic data that was collected from other OS-based monitors.

Do I still need my OS-based performance tools?

PowerVP was created to complement your current suite of performance tools in your toolbox.

PowerVP focuses on new views that typically are not available with your OS-based performance tools. Use these tools together to monitor and optimize your system/application performance.

How should I set the color coded CPU utilization thresholds?

The utilization of a particular resource (core, disk drive, bus) simply indicates its level of being busy doing work. High or low doesn't necessarily imply good or bad. If there is important work offered to your Power System, you would hope that it would be run now; and processing that work drives higher resource utilization. If there is spare resource utilization, you might want that low-priority batch jobs are able to take advantage of it; and processing that work drives higher resource utilization.

If there is spare resource utilization, you might want to be energy conscious and have some cores put to sleep; the result of this action drives the utilization of the remaining resources to a higher utilization. The point here is that high utilization isn't necessarily a bad thing.

It is important to do sizing and capacity planning such that your system resources can handle the anticipated load as well as reasonable peaks in workload. In this planning, it is also important to plan some headroom (for example, spare utilization). Part of the purpose of headroom, is to be able to do most of your work with CPU utilization level low enough not to have too much queuing time. The other purpose of headroom is to be able to handle workload peaks; perhaps during some of those peaks, you can deal with having additional response time due to the queuing multiplier effect. Best practices for headroom levels consider many factors (number of cores, resource type, partition type, partition size, and so on). Use the IBM Systems Workload Estimator tool to size a new system, to size an upgrade to an existing system, or to size a consolidation of several systems. The IBM Systems Workload Estimator tool can be found at www.ibm.com/systems/support/tools/estimator.

Therefore, you should customize these color-coded thresholds to meet the requirements of your business. You can set the number of thresholds, the utilization levels, and the colors. You might start with the default levels/colors and modify them to your customized environment. This customization helps to set your expectations and actions as to what to do when those levels are exceeded. For example if you see that your CPU utilization is generally red for an hour during your business day, do you: 1) quickly change to increase the entitlement of a high priority LPAR, 2) consider a hardware upgrade to migrate to a newer/larger Power System soon, 3) consider activating some additional capacity on demand, or 4) realize that your system experienced a normal peak in a red zone of an average day?

If I see red, do I have a performance problem?

Probably not, reread the previous section. Red indicates a high utilization for a resource.

How should I set the color coded link utilization thresholds?

The availability of the instrumentation for the Powerbus, MC bus, and GX/PHB bus is relatively recent. The default utilization thresholds/colors are set as a starting point. The default thresholds are monitored and improved by using various workloads within IBM and from feedback from customers who use PowerVP. High or imbalanced Powerbus utilization can be an indication that improvements can be made to affinity (see the affinity section).

How should I monitor my system with PowerVP?

It depends on the nature of your business and the health of your servers. You may wish to always record data and monitor it in real time, and then use the play back functions to do drill-down analysis as needed. To drill deeper, simply use the PowerVP navigation (clicking and hovering) to look more closely at hardware nodes, bus utilizations, and partition details. Remember you can monitor only (whether realtime or with play-back) information that was recorded with PowerVP as it cannot use historic data that is collected from other OS-based monitors. It is also possible to record PowerVP data without running the monitor GUI. Some clients have talked about having the PowerVP system-level display projected on their wall or on the "big screen". They can customize PowerVP to alert them with particular colors for certain utilization levels.

How can I try optimization ideas?

You want to obtain a good “before” monitoring interval prior to any changing of application or configuration. In doing so, note the date and time for the replay, or take screen shots, or note the CPU utilizations and CPI levels. For the LPAR drill-down panels, you can mark the bar charts with blue marks that indicate the current levels. Then do the change to your application or configuration that you assume provides an optimization. After this change settles down look again at the PowerVP data to confirm the performance level after the change. Typically one would want to keep the workload level equivalent to make appropriate comparisons. Now you can look for improvement indicators: CPU utilization reduction, CPI reduction, movement of the LSU CPI breakdown from right to left (remote memory to local memory, remote caches to local caches, L3 to L2, and so on), reduction of the utilizations of buses (Powerbus, MC bus, GX bus).

How can PowerVP help with DPO?

The Dynamic Platform Optimizer (DPO) can help optimize your virtualization configuration. Just like making any other change on your own, you can use PowerVP to help validate the performance benefit of DPO.

How do I map virtual partitions to physical configurations?

The new hypervisor interfaces created for PowerVP provide topology information. The illustrations on the PowerVP monitor show the specific existence and topology of cores, chips, processor modules, hardware nodes, and links. Each major view also has an LPAR section to show the virtual perspective by listing the partitions. PowerVP is able to help you map the virtual partitions to the physical configuration. By clicking a dedicated partition, PowerVP highlights that partition in a unique color and also lights up the CPU resources (cores) in the same color. Then you can look at this mapping. Ideally, from the perspective of a given partition, the cores that are allocated would be grouped close together in the configuration to maximize the locality of data access. For shared processor partitions, PowerVP will map a given partition to a shared processor pool. Tasks in shared partitions can be dispatched to any core in the shared processor pool. If your server configuration has recently changed by adding partitions or changing the entitlement, you may want to consider running DPO to optimize the configuration for performance.

How can I tune for better affinity?

With IBM Power Systems, having good affinity is important for good performance. Servers that use a nodal design to increase their capacity (e.g., NUMA-like architectures) are especially sensitive to affinity considerations. Processor affinity suggests that your work be dispatched to hardware threads to the cores/chips/nodes with the highest likelihood of being close in proximity where your data is. Memory affinity suggests that memory that is allocated to your work be close in proximity to the cores processing your work. Ideally, your work would be dispatched to the same core to optimize the chance of having a hot cache (vs. a dirty cache, or having to do remote cache accesses) or at least to the same socket or node to optimize the chance of having local memory accesses (vs. having to do remote/distant memory accesses). Much of the CPU resource that is consumed for your application can be attributed to accessing data (that is, consuming cycles while waiting for accesses from cache or memory). When accessing data, the goal is to consume as few cycles as possible. For POWER7 and POWER8, it is preferred to access data from this list in this order of preference: L1 cache, local L2 cache, local L3 cache, local L4 cache, cache on another core on the same chip, cache on another chip on the same hardware node, cache on another hardware node, memory on your socket, memory on another socket on the same hardware node, memory on another hardware node. There might be a 1000 times difference in run cycles that are consumed from one extreme to another. To improve your affinity, you could try a number of things. Keep in mind that this can be an advanced technical topic. Application coding adjustments may maximize cache line optimization. Using pre-fetch or not may provide trade-offs between CPI, throughput, and response time. Using virtualization, such as dedicated partitions, may force better affinity. Using other OS-provided functions (RSET, subsystems, WPAR, affinity system values, etc.) may force better affinity. Many of these topics are discussed in papers at: www.ibm.com/systems/power/software/i/management/performance/resources.html.

Note: Local L4 cache is only available on POWER8.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licenseses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

Programming interface information

This Logical partitioning publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM AIX Version 7.1, IBM AIX Version 6.1, IBM i 7.2, and IBM Virtual I/O Server Version 2.2.4.00.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Printed in USA