

IBM PowerVM for Growing Businesses: Managing and Monitoring a Virtual Environment

IBM Redbooks Solution Guide

This IBM® Redbooks® Solution Guide introduces IBM PowerVM® virtualization technologies on IBM Power Systems™ servers. In a virtualized environment, and with the intrinsic capabilities of PowerVM, systems administrators can perform management and monitoring tasks more efficiently. PowerVM provides a secure and scalable virtualization environment for IBM AIX®, IBM i, and Linux applications. This solution guide describes several PowerVM features that can increase management and monitoring efficiencies. In particular, it offers guidance for planning and setting up the multiple shared processor pool feature.

A companion solution guide (TIPS1099; see “Related information”) expands on these ideas and delves into the decreased total cost of computing made possible by PowerVM, along with additional management and monitoring techniques. This solution guide is for you if you have the drive to maximize the efficiency of management and monitoring activities in a virtual environment. Figure 1 presents the IBM Power Systems family of products.



Figure 1. PowerVM provides virtualization technologies for the IBM Power Systems family of products

Did you know?

In a virtualized environment, the operating system sees only the virtual processors and dispatches runnable tasks to them in the normal course of running a workload. The focus of this solution guide is to introduce you to the processor, memory, and I/O server virtualization features of PowerVM, and to describe the efficiencies offered by PowerVM virtualization, particularly in the areas of systems management and monitoring.

Business value

Maximize the return on your IT infrastructure investments by consolidating workloads. IBM Power Systems, combined with PowerVM technology, is designed to help consolidate and simplify your environment. Key capabilities include the following benefits:

- Improve server utilization and I/O resource sharing; reduce cost of ownership; use IT assets wisely.
- Improve business responsiveness and operational speed. Dynamically reallocate application resources, as needed, to meet changing business needs or handle unexpected changes in demand.
- Simplify IT infrastructure management by making workloads independent of hardware resources. Make business-driven policies to deliver resources based on time, cost, and service-levels.

What is PowerVM

PowerVM provides the industrial-strength virtualization solution for IBM Power Systems servers and blades. This solution provides proven workload consolidation that helps clients control costs and improves overall performance, availability, flexibility and energy efficiency. PowerVM is a combination of hardware enablement and value-added software. The primary features and technologies (and function provided by) of PowerVM are indicated in the following list. These features are described in detail in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 (see "Related information").

- PowerVM Hypervisor (hardware platform)
- Logical partitioning (hypervisor)
- Micro-partitioning (hypervisor)
- Dynamic logical partitioning (hypervisor)
- Shared processor pools (hypervisor)
- Integrated Virtualization Manager (hypervisor, IVO, IVM)
- Shared storage pools (hypervisor, VIOS)
- Virtual I/O Server (hypervisor, VIOS)
- Virtual SCSI (hypervisor, VIOS)
- Virtual Fibre Channel (hypervisor, VIOS)
- Virtual optical device and tape (hypervisor, VIOS)
- Live partition mobility (hypervisor, VIOS)
- Partition suspend/resume (hypervisor, VIOS)
- IBM Active Memory™ Sharing (hypervisor, VIOS)
- Active Memory Deduplication (hypervisor)
- Active Memory Mirroring (hypervisor)
- Host Ethernet Adapter, or HEA (hypervisor)

The following technologies are also frequently mentioned together with PowerVM. They are detailed in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 (see "Related information").

- IBM POWER® processor compatibility modes (hypervisor)
- Capacity on Demand (hypervisor)
- Simultaneous Multithreading (hardware, AIX)
- Active Memory Expansion (hardware, AIX); available only on IBM POWER7® Systems and later
- AIX Workload Partitions (AIX); available only on AIX version 6.1 or later
- System Planning Tool, known as SPT (SPT)

PowerVM editions

Three versions of PowerVM are available, each suited to various purposes:

- *PowerVM Express Edition:* PowerVM Express Edition is designed for customers looking for an introduction to more advanced virtualization features at a highly affordable price.
- *PowerVM Standard Edition:* PowerVM Standard Edition provides the most complete virtualization functionality for AIX, IBM i, and Linux operating systems in the industry. PowerVM Standard Edition is supported on Power Systems servers and includes features designed to allow businesses to increase system utilization.
- *PowerVM Enterprise Edition:* PowerVM Enterprise Edition includes all the features of PowerVM Standard Edition plus two new industry-leading capabilities called Active Memory Sharing and Live Partition Mobility.

Upgrading is possible: from Express Edition to either Standard or Enterprise Edition, and from Standard Edition to Enterprise Edition. Table 3 lists functional elements of the available PowerVM editions.

Table 3. Overview of PowerVM capabilities by edition

PowerVM capability	PowerVM Express Edition	PowerVM Standard Edition	PowerVM Enterprise Edition
Maximum VMs	3 / Server	1000 / Server	1000 / Server
Micro-partitions*	Yes	Yes	Yes
Virtual I/O Server	Yes (single)	Yes (dual)	Yes (dual)
Management	VMControl, IVM	VMControl, IVM**, HMC	VMControl, IVMb, HMC
Shared dedicated capacity	Yes	Yes	Yes
Multiple shared processor pools +	No	Yes	Yes
Live Partition Mobility	No	No	Yes
Active Memory Sharing++	No	No	Yes
Active Memory++	No	No	Yes
Suspend/Resume	No	Yes	Yes
Virtual Fibre Channel	Yes	Yes	Yes
Shared Storage Pools	No	Yes	Yes
Thin provisioning	No	Yes	Yes
Thick provisioning	No	Yes	Yes

* When firmware is at level 7.6 or later, micro-partitions can be defined as small as 0.05 of a processor instead of 0.1 of a processor.

** IVM supports only a single Virtual I/O Server.

+ Needs IBM POWER6 processor-based system or later.

++ Needs IBM POWER7 processor-based system with firmware at level 7.4 or later.

Read an overview of the availability of the PowerVM features by Power Systems models:

<http://www.ibm.com/systems/power/software/virtualization/editions/features.html>

The PowerVM feature is a combination of hardware enablement and software that are available together as a single priced feature. It is charged at one unit for each activated processor, including software maintenance. In addition, not the following information:

- The software maintenance can be ordered for a one-year or three-year period. It is also charged for each active processor on the server.
- When the hardware feature is specified with the initial system order, the firmware is shipped already activated to support the PowerVM features.
- For an HMC-attached system with the PowerVM Standard Edition or the PowerVM Enterprise Edition, the processor-based license enables you to install several Virtual I/O Server partitions (usually two) on a single physical server to provide redundancy and to spread the I/O workload across several Virtual I/O Server partitions.
- Virtual Ethernet and dedicated processor partition are features available without the PowerVM feature for servers attached to an HMC.

The IBM POWER Hypervisor

The IBM POWER Hypervisor™ is the foundation of IBM PowerVM. The POWER Hypervisor provides the ability to divide physical system resources into isolated logical partitions. Each logical partition operates like an independent system running its own operating environment: AIX, IBM i, Linux, or the Virtual I/O Server. The Hypervisor can assign dedicated processors, I/O, and memory, which you can dynamically reconfigure as needed, to each logical partition.

The Hypervisor can also assign shared processors to each logical partition using its micro-partitioning feature. Unknown to the logical partitions, the Hypervisor creates a shared processor pool from which it allocates virtual processors to the logical partitions as needed. This means the Hypervisor creates virtual processors so that logical partitions can share the physical processors while running independent operating environments.

Combined with features designed into the IBM POWER processors, the POWER Hypervisor delivers functions that enable capabilities including dedicated-processor partitions, micro-partitioning, virtual processors, IEEE VLA- compatible virtual switch, virtual Ethernet adapters, virtual SCSI adapters, virtual Fibre Channel adapters, and virtual consoles.

The POWER Hypervisor is a firmware layer sitting between the hosted operating systems and the server hardware, as shown in Figure 2.

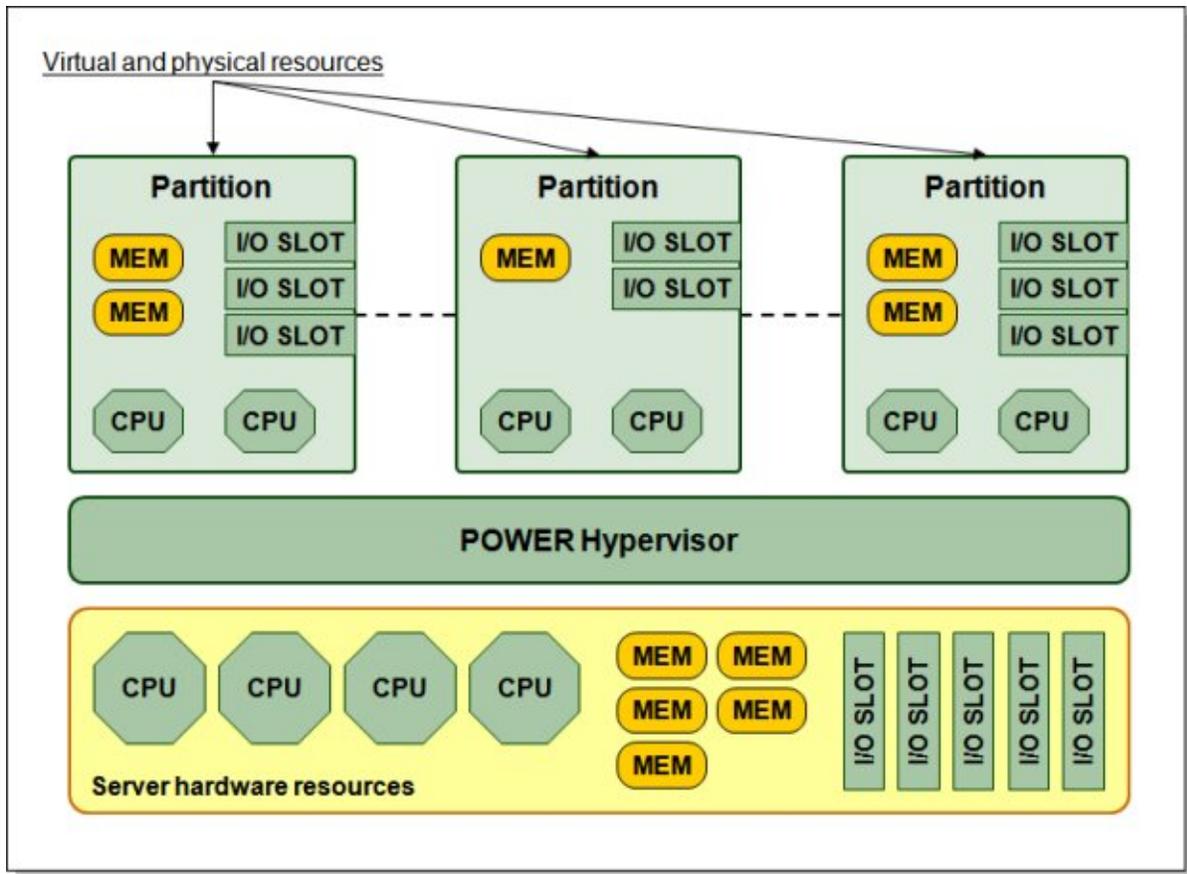


Figure 2. POWER Hypervisor abstracts physical server hardware

I/O virtualization

Combined with features designed into the POWER processors, the POWER Hypervisor delivers functions that enable other system technologies, including logical partitioning technology, virtualized processors, IEEE VLAN compatible virtual switch, virtual SCSI adapters, virtual Fibre Channel adapters and virtual consoles. The POWER Hypervisor is a basic component of the system's firmware and offers the following functions:

- Provides an abstraction between the physical hardware resources and the logical partitions that use them
- Enforces partition integrity by providing a security layer between logical partitions
- Controls the dispatch of virtual processors to physical processors.
- Saves and restores all processor state information during a logical processor context switch
- Controls hardware I/O interrupt management facilities for logical partitions
- Provides virtual LAN channels between logical partitions that help to reduce the need for physical Ethernet adapters for inter-partition communication
- Monitors the Service Processor and will perform a reset/reload if it detects the loss of the Service Processor, notifying the operating system if the problem is not corrected

The POWER Hypervisor is always active, regardless of the system configuration and also when not connected to the HMC. The POWER Hypervisor provides the following types of virtual I/O adapters:

- Virtual SCSI
- Virtual Ethernet
- Virtual Fibre Channel
- Virtual console

Virtual I/O Server overview

As part of PowerVM, the Virtual I/O Server is a software appliance with which you can associate physical resources and that allows you to share these resources among multiple client logical partitions. The Virtual I/O Server can provide both virtualized storage and network adapters, making use of the virtual SCSI and virtual Ethernet facilities. Figure 3 shows a very basic Virtual I/O Server configuration. This diagram shows only a small subset of the capabilities to illustrate the basic concept of how the Virtual I/O Server works. As shown, the physical resources such as the physical Ethernet adapter and the physical disk adapter are accessed by the client partition using virtual I/O devices.

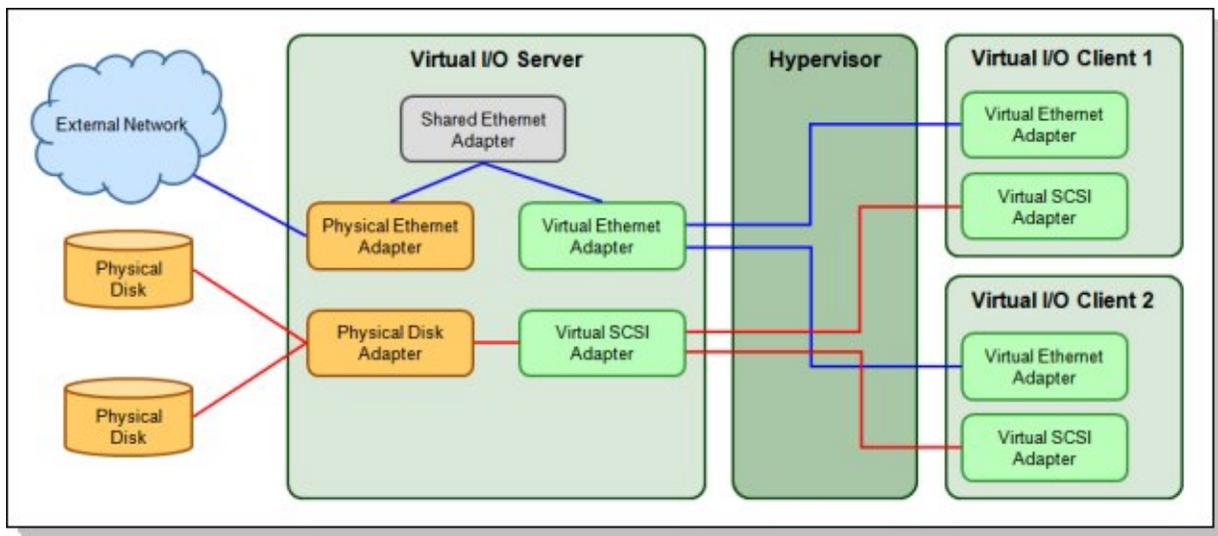


Figure 3. Simple Virtual I/O Server configuration

Solution overview

This section outlines the overall capabilities of PowerVM, then focuses on the capabilities of processor virtualization specifically.

PowerVM virtualization capabilities

The primary virtualization capabilities that PowerVM offers are as follows:

- *Processor virtualization.* The virtualization of physical processors in IBM Power Systems introduces an abstraction layer that is implemented within the IBM POWER Hypervisor. The POWER Hypervisor abstracts the physical processors and presents a set of virtual processors to the operating system within the micro-partitions on the system.
- *Memory virtualization.* POWER technology-based servers are very powerful and provide much processor capacity. Memory is therefore often the bottleneck that prevents an increase in the overall server utilization.

- *I/O virtualization:* Combined with features designed into the POWER processors, the POWER Hypervisor delivers functions that enable other system technologies, including logical partitioning technology, virtualized processors, IEEE VLAN compatible virtual switch, virtual SCSI adapters, virtual Fibre Channel adapters and virtual consoles. The POWER Hypervisor is a basic component of the system's firmware and offers the following functions:
 - Provides an abstraction between the physical hardware resources and the logical partitions that use them.
 - Enforces partition integrity by providing a security layer between logical partitions.
 - Controls the dispatch of virtual processors to physical processors. Saves and restores all processor state information during a logical processor context switch.
 - Controls hardware I/O interrupt management facilities for logical partitions.
- *Server virtualization:* These features provide the following capabilities:
 - Live partition mobility: PowerVM Live Partition Mobility allows you to move a running logical partition, including its operating system and running applications, from one system to another without any shutdown or without disrupting the operation of that logical partition. Inactive partition mobility allows you to move a powered off logical partition from one system to another.
 - Partition suspend and resume: The Virtual I/O Server provides partition suspend and resume capability to client logical partitions within the IBM POWER7 systems. With suspend and resume operations, the partition's state can be suspended, and then resumed at a later time.

Processor virtualization capabilities

Technologies associated with processor virtualization on PowerVM are as follows:

- *Micro-partitioning:* This is the ability to distribute the processing capacity of one or more physical processors among one or more logical partitions. Thus, processors are shared among logical partitions. Micro-partitioning is supported across the entire POWER5 and later server range, from entry level to the high-end systems. The benefit of micro-partitioning is that it allows significantly increased overall utilization of processor resources within the system. The micro-partition is provided with a processor entitlement, which is the processor capacity guaranteed to it by the POWER Hypervisor.
- *Shared processor pools:* Micro-partitioning technology coupled with the POWER Hypervisor facilitates the sharing of processing units between micro-partitions. Processor resources can be used very efficiently with Shared Process Pools (SPP), leading to a significantly increased overall system utilization. A shared processor pool is a specific group of micro-partitions (and their associated virtual processors) through which processor capacity from the physical shared processor pool can be managed.
- *Multiple shared processor pools:* Multiple shared processor pools (MSPPs) is a capability supported on POWER6 (or later) technology. This capability allows the user to define shared processor pools with the purpose of controlling the processor capacity that can be consumed from the physical shared processor pool. Multiple shared processor pools can be used to isolate workloads in a pool and thus not exceed an upper CPU limit. Multiple shared processor pools can also be useful for software license management where sub-capacity licensing is involved.
- *Shared dedicated capacity:* POWER6-based and later servers offer the capability of harvesting unused processor cycles from dedicated-processor partitions. These unused cycles are then donated to the physical shared processor pool associated with micro-partitioning. This ensures the opportunity for maximum processor utilization throughout the system.

With these technologies defined, we plan for and set up the multiple shared processor pools.

Solution architecture

All Power Systems servers that support the multiple shared processor pools capability will have a minimum of one (the default) shared processor pool and up to a maximum of 64 shared processor pools. Micro-partitions are created and then identified as members of either the default shared processor pool (SPP0) or a user-defined shared processor pool (SPPn). The virtual processors that exist within the set of micro-partitions are monitored by the POWER Hypervisor and processor capacity is managed according to user-defined attributes.

An overview of the architecture of multiple shared processor pools can be seen in Figure 4.

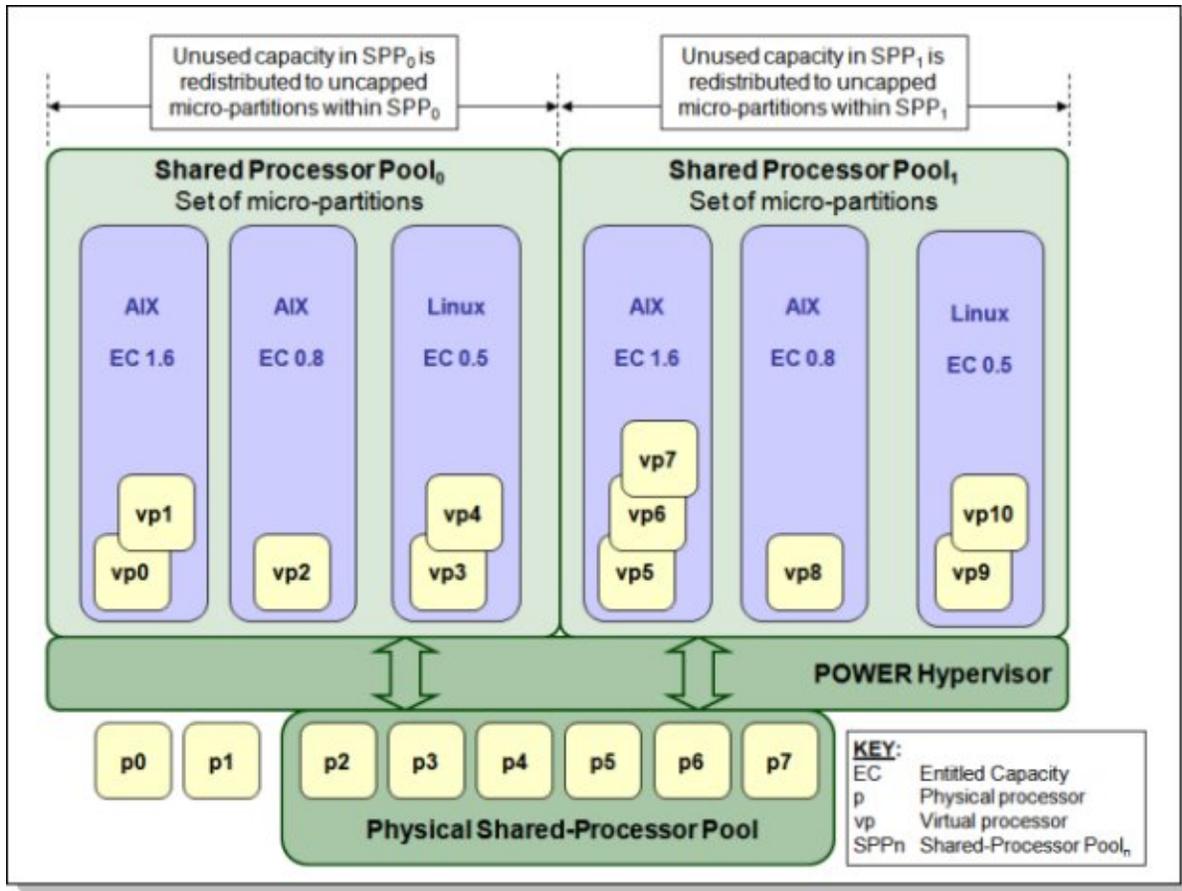


Figure 4. Overview of the architecture of multiple shared processor pools

The diagram shows a system with eight processors (p0 to p7) and two shared processor pools (the default SPP0 and the user-created SPP1), with three micro-partitions each. Each micro-partition is seen by operating system as an individual system and may run a different AIX or Linux version.

Processors p0 and p1 are not donating capacity and, consequently, are not part of the physical shared processor pool. The total capacity of processors p2 to p7 is distributed among the eleven virtual processors (vp0 to vp10) available on the micro-partitions. The POWER Hypervisor guarantees the minimum entitled capacity of each partition. Unused capacity in SPP0 and SPP1 is redistributed among their micro-partitions.

For more details about micro-partitioning and SPP capacity concepts, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 (see "Related information").

Live Partition Mobility and multiple shared processor pools

A micro-partition can leave a shared processor pool because of PowerVM Live Partition Mobility. Similarly, a micro-partition can join a shared processor pool in the same way. When performing PowerVM Live Partition Mobility, you are given the opportunity to designate a destination shared processor pool on the target server to receive and host the migrating micro-partition.

Because several simultaneous micro-partition migrations are supported by PowerVM Live Partition Mobility, migrating the entire shared processor pool from one server to another is conceivable.

Shared dedicated capacity

POWER6 technology-based and later servers offer the capability of harvesting unused processor cycles from dedicated-processor partitions. These unused cycles are then donated to the physical shared processor pool associated with micro-partitioning. This ensures the opportunity for maximum processor utilization throughout the system.

Usage scenarios

In one scenario, a client wants to use the multiple shared processor pool feature of PowerVM. Begin with the planning phase. A second scenario is also described.

First scenario: Shared processor pools

The planning phase for setting up multiple shared processor pools includes an evaluation of the capacity attributes. Here, we describe these capacity attributes and provide examples of the capacity resolution according to the server load.

Capacity attributes

The following attributes are used for calculating the pool capacity of shared processor pools:

- **Maximum Pool Capacity:** Each shared processor pool has a maximum capacity associated with it. The Maximum Pool Capacity (MPC) defines the upper boundary of the processor capacity that can be utilized by the set of micro-partitions in the shared processor pool. The Maximum Pool Capacity must be represented by a whole number of processor units.
- **Reserved Pool Capacity:** The system administrator can assign an entitled capacity to a shared processor pool for the purpose of reserving processor capacity from the physical shared processor pool for the express use of the micro-partitions in the shared processor pool. The Reserved Pool Capacity (RPC) is in addition to the processor capacity entitlements of the individual micro-partitions in the shared processor pool. The Reserved Pool Capacity is distributed among uncapped micro-partitions in the shared processor pool according to their uncapped weighting. Default value for the Reserved Pool Capacity is zero (0).
- **Entitled Pool Capacity:** The Entitled Pool Capacity (EPC) of a shared processor pool defines the guaranteed processor capacity that is available to the group of micro-partitions in the shared processor pool. The Entitled Pool Capacity is the sum of the entitlement capacities of the micro-partitions in the shared processor pool plus the Reserved Pool Capacity.

Using the information in Table 4 as an example and a Reserved Pool Capacity of 1.5, calculating the Entitled Pool Capacity is easy.

Table 4. Entitled capacities for micro-partitions in a shared processor pool

Micro-partitions in a shared processor pool	Entitled capacity for micro-partitions
Micro-partition 0	0.5
Micro-partition 1	1.75
Micro-partition 2	0.25
Micro-partition 3	0.25
Micro-partition 4	1.25
Micro-partition 5	0.50

The sum of the entitled capacities for the micro-partitions in this shared processor pool is 4.50, which gives us the following calculation:

$$\text{EntitledPoolCapacity} = 4.50 + 1.50$$

This gives the following result:

$$\text{EntitledPoolCapacity} = 6.0$$

Setup phase

The setup phase for multiple shared processor pools includes the following items:

- Processor virtualization setup
- Configuring multiple shared processor pools
- Dynamic adjustment of maximum pool capacity
- Dynamic adjustment of reserve pool capacity
- Dynamic movement between shared processor pools

Processor virtualization setup

This section shows the basics of setting up multiple shared processor pools and shared-dedicated capacity. Now that you are familiar with the concepts and terms described in *Process virtualization capabilities*, read about how to set up multiple shared processor pools and shared-dedicated capacity.

For advanced setup and maintenance information, see *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590 (see "Related information").

Configuring multiple shared processor pools

All Power Systems that support the multiple shared processor pools capability will have a minimum of one (the default) shared processor pool and up to a maximum of 64 shared processor pools.

The default shared processor pool (SPP0) is automatically activated by the system and is always present. Its Maximum Pool Capacity is set to the capacity of the physical shared processor pool. For SPP0, the Reserved Pool Capacity is always 0.

All other shared processor pools exist, but by default, are inactive. By changing the Maximum Pool Capacity of a shared processor pool to a value greater than zero, it becomes active and can accept micro-partitions (either transferred from SPP0 or newly created).

The system administrator can use the HMC to activate additional multiple shared processor pools. As a minimum, the Maximum Pool Capacity will need to be specified. If you want to specify a Reserved Pool Capacity, there must be enough unallocated physical processor capacity to guarantee the entitlement.

Dynamic adjustment of Maximum Pool Capacity

The Maximum Pool Capacity of a shared processor pool, other than the default SPP0, can be adjusted dynamically from the HMC using either the graphical or command-line interface (CLI).

Dynamic adjustment of Reserve Pool Capacity

The Reserved Pool Capacity of a shared processor pool, other than the default SPP0, can be adjusted dynamically from the HMC using either the graphical interface or CLI.

Dynamic movement between shared processor pools

A micro-partition can be moved dynamically from one shared processor pool to another from the HMC using either the graphical or CLI interface. As the Entitled Pool Capacity is partly made up of the sum of the entitled capacities of the micro-partitions, removing a micro-partition from a shared processor pool will reduce the Entitled Pool Capacity for that shared processor pool. Similarly, the Entitled Pool Capacity of the shared processor pool that the micro-partition joins will increase.

Second scenario: Configuring multiple shared-processor pools to limit processing capacity

This scenario is about creating multiple shared processor pools to limit the amount of processing capacity that can be consumed by the micro-partitions, based on the micro-partitions entitlement, virtual processors, pool name, and pool ID, as described in Table 5. In this scenario, all micro-partitions are running in uncapped mode.

Table 5. Micro-partition configuration and shared processor pool assignments

Partition name	Entitled Capacity	Number of virtual processors	Pool name	Pool ID
VIO_Server1	0.3	2	DefaultPool	0
VIO_Server2	0.3	2	DefaultPool	0
DB_server	1	3	Production	1
Apps_server	1	3	Production	1
NIM_server	0.2	1	Development	2
linuxlpar	0.2	1	Development	2

The two Virtual I/O Servers are created in the default shared processor pool (Pool ID = 0). On the HMC this pool is referred to as DefaultPool. DefaultPool cannot be modified. Each of the Virtual I/O Servers has an entitlement of 0.3 processing units and 2 virtual processors configured. Each one is therefore guaranteed to get 0.3 processing units and can consume up to 2 processing units if required and if spare processing capacity is available in the DefaultPool.

The DB_server and the Apps_server micro-partitions are assigned to the shared processor pool called Production (Pool ID = 1). As shown in Table 6, the Maximum Pool Capacity of the Production pool is limited to 3. A reason for limiting a production workload might be to reduce licensing cost for applications running in these partitions.

Table 6. Shared processor pool attributes

Pool name	Maximum pool capacity	Entitled pool capacity	Reserved pool capacity
Production	3	2.5	0.5
Development	1	0.4	0

The Entitled Pool Capacity, which is the amount of processing units that the pool is guaranteed to get, is 2.5 for the Production pool. This is the sum of the entitlements of the partitions in the pool (in this case 1.0 for the DB_server and 1.0 for the Apps_server) plus the Reserved Pool Capacity of 0.5 processing units. The Reserved Pool Capacity is distributed within the bounds of the pool according to the weighting factors defined for each partition. A reason for using Reserved Pool Capacity might be to guarantee unallocated processor resource to balance workload demand in production workloads.

Although the DB_server and the Apps_server micro-partitions each have three virtual processors configured, the total processing capacity, which can be used by the two micro-partitions at a time, is limited to 3 physical processors. This is because the Maximum Pool Capacity in the shared processor pool. If they were assigned to the DefaultPool, they can use up to 6 physical processors if capacity is available.

The NIM_server and the linuxlpar partitions are assigned to the Development pool. The Development pool has a maximum pool capacity of 1 processing unit. The micro-partitions in the Development pool can use at most the capacity of 1 physical processor. The Entitled Pool Capacity of the Development pool is 0.4 as shown in Table 6.

Tip: The pool attributes and partition assignments can be changed dynamically while the micro-partitions are running.

For additional tasks, including the following, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 (see "Related information"):

- Shared processor pool management using the HMC GUI
- Shared processor pool management using the command line

Shared dedicated capacity

The system administrator can control which dedicated-processor partitions can donate unused cycles. The dedicated-processor partition must be identified as a donating partition.

When the CPU utilization of the core goes below a threshold, and all the SMT threads of the CPU idle from a hypervisor perspective, the CPU will be donated to the shared processor pool.

The operating system will make a thread idle from hypervisor perspective when it enters the idle loop and the SMT snooze delay expires. The delay needs to be set to zero to maximize the probability of donation. Other than the SMT snooze delay, the donation completes in microseconds.

The process for controlling which dedicated-processor partitions can donate unused cycles is described in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 (see "Related information"). The publication includes information pertaining to the AIX, IBM i, and Linux operating systems.

The donated processor is returned instantaneously (within microseconds) to the dedicated processor partition when one of the following conditions occurs:

- The timer of one of the SMT threads on the donated CPU expires.
- An external interrupt for the dedicated-processor partition is presented to one of the SMT threads of the donated CPU.
- The operating system needs the CPU back to dispatch work on one of the SMT threads of the donated CPU.

LPARs that use processor folding will tend to donate more idle capacity because the workload is constrained to a subset of the available processors and the remaining processors are ceded to the hypervisor on a longer term basis than they are when snooze-delay techniques are used.

A workload in a shared-dedicated partition might see a slight performance impact because of the cache-effects of running micro-partitions on a donated CPU.

Switching a partition from shared-dedicated to dedicated or reverse is a dynamic LPAR operation.

Instructions for setting up a dedicated donating processor on HMC are provided in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 (see "Related information").

Integration

Integration with the following components is described in this section:

- Hardware Management Console (HMC)
- Integrated Virtualization Manager (IVM)
- System Director VMControl

Hardware Management Console

The Hardware Management Console (HMC) is a dedicated Linux-based appliance that you use to configure and manage IBM Power System servers. The HMC provides access to logical partitioning functions, service functions, and various system management functions through both a browser-based interface and a command line interface (CLI). Because it is a separate stand-alone system, the HMC does not use any managed system resources and you can maintain it without affecting system activity.

IBM PowerVM technology and HMC has been enhanced with the following features:

- Support up to 16 concurrent Live Partition Mobility (LPM) activities.
- DLPAR Add or Remove of virtual I/O adapters to or from a Virtual I/O Server (VIOS).

HMC now automatically attempts to run the Add/Remove commands (`cfgdev/rmdev`) on the VIOS for the user. Prior to this enhancement, the user had to manually run these commands on the VIOS.

Integrated Virtualization Manager

The Integrated Virtualization Manager (IVM) is used to manage selected Power Systems servers using a web-based graphical interface without requiring an HMC.

This reduces the hardware needed for the adoption of virtualization technology, particularly for low-end systems. This solution fits in small and functionally simple environments where only few servers are deployed or some advanced HMC-like functions are required.

The Integrated Virtualization Manager (IVM) is a basic hardware management solution, included in the VIO software that inherits key Hardware Management Console (HMC) features.

System Director VMControl

IBM Systems Director is a platform-management foundation that streamlines the way you manage physical and virtual systems across a heterogeneous environment. By using industry standards, IBM System Director supports multiple operating systems and virtualization technologies across IBM and non-IBM x86 platforms. IBM Systems Director has three editions: Express, Standard, and Enterprise Editions.

IBM Systems Director VMControl is a plug-in for IBM Systems Director. It is a cross-platform suite of product that assists you in rapidly deploying virtual appliances to create virtual servers that are configured with the operating system and software applications that you want. It also enables you to group resources into system pools, which enables you to centrally manage and control the workloads in your environment.

The IBM System Director VMControl plug-in is available in four editions: VMControl Express Edition, VMControl Standard Edition, VMControl Enterprise Edition, and IBM System Director VMControl for IBM PowerLinux™, each with an increasing set of functionality for virtualization management. Detailed descriptions of these editions are included in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 (see "Related information").

Supported platforms

PowerVM supports the following operating systems and operating components:

- Version 7, Release 1 of AIX
- Version 7, Release 1 of IBM i
- Version 2, Release 2, Modification 2, Fixpack 26 of the Virtual I/O Server
- Version 7, Release 7, Modification 6 of the HMC
- Version AL730, release 95 of the POWER7 System Firmware
- Version AL740, release 95 of the POWER7 System Firmware

Hardware requirements

The Virtual I/O Server can run on any POWER5 or later server that has the PowerVM Standard feature enabled. Also supported are IBM BladeCenter® Power Blade servers. With the PowerVM Standard or PowerVM Enterprise Edition, Virtual I/O Servers can be deployed in pairs to provide high availability.

PowerVM features are supported on the majority of the Power Systems offerings, however, there are some exceptions. The *Availability of PowerVM features by Power Systems models* web page has a summary of server models and features at:

<http://www.ibm.com/systems/power/software/virtualization/editions/features.html>.

Table 7 lists the Power technology levels that provide individual features.

Table 7. Virtualization features supported by POWER technology levels

Feature	POWER5	POWER6	POWER7
Virtual SCSI	Yes	Yes	Yes
Virtual Ethernet	Yes	Yes	Yes
Shared Ethernet Adapter	Yes	Yes	Yes
Virtual Fibre Channel	No	Yes	Yes
Virtual Tape	Yes	Yes	Yes
Logical partitioning	Yes	Yes	Yes
DLPAR I/O adapter add/remove	Yes	Yes	Yes
DLPAR processor add/remove	Yes	Yes	Yes
DLPAR memory add/remove	Yes	Yes	Yes
Micro-partitioning	Yes	Yes	Yes
Shared Dedicated Capacity	Yes*	Yes	Yes
Multiple shared processor pools	No	Yes	Yes
Virtual I/O Server	Yes	Yes	Yes
Integrated Virtualization Manager	Yes	Yes	Yes
Suspend and resume	No	No	Yes
Shared storage pools	No	Yes	Yes
Thin provisioning	No	Yes	Yes
Active Memory Sharing	No	Yes	Yes
Active Memory Deduplication	No	No	Yes**
Active Memory Mirroring	No	No	Yes+
Live Partition Mobility	No	Yes	Yes

Table 7. Virtualization features supported by POWER technology levels (continued)

Feature	POWER5	POWER6	POWER7
Simultaneous multithreading	Yes++	Yes	Yes+++
Active Memory Expansion	No	No	Yes
Capacity on Demand+++	Yes	Yes	Yes
AIX Workload Partitions	Yes	Yes	Yes

* Only capacity from shutdown partitions can be shared

** Need firmware level 7.4, or later

+ Need mid-tier and large-tier IBM POWER7 Systems™ or later, including Power 770, 780, and 795

++ POWER5 supports 2 threads

+++ POWER7 or later supports 4 threads

Licensing requirements

Before using PowerVM, you also need to have a valid license, or we can say feature code. You need to check your order and confirm that the feature code is included. If you order an upgrade feature, you need to input and activate the code first. The procedure to activate the code is described in *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940 (see "Related information").

Ordering information

Ordering information is shown in the following table.

Table 8. PowerVM feature codes

Type and model of IBM Power System	Express Edition	Standard Edition*	Enterprise Edition**
7895-22X	5225	5227	5228
7895-23X	5225	5227	5228
7895-42X	5225	5227	5228
8202-E4B	5225	5227	5228
8202-E4C	5225	5227	5228
8205-E6C	5225	5227	5228
8205-E6B	5225	5227	5228
8231-E2B	5225	5227	5228
8231-E1C and 8231-E2C	5225	5227	5228
8233-E8B	7793	7794	7795
9117-MMB	Not offered	7942	7995
9117-MMC	Not offered	7942	7995
9117-MMD	Not offered	7942	7995
9119-FHB	Not offered	7943	8002
9179-MHB	Not offered	7942	7995
9179-MHC	Not offered	7942	7995
9179-MHD	Not offered	7942	7995

* The feature codes for the Standard Edition provide all functions supplied with the Express Edition.

** The feature codes for the Enterprise Edition provide all functions supplied with the Standard Edition.

Related information

For more information, see the following documents:

- IBM Offering Information page (to search on announcement letters, sales manuals, or both):
http://www.ibm.com/common/ssi/index.wss?request_locale=en
On this page, enter `IBM PowerVM`, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.
- *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
<http://www.redbooks.ibm.com/abstracts/sg247940.html>
- *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590
<http://w3.itso.ibm.com/abstracts/sg247590.html?Open>
- *IBM PowerVM for Growing Businesses: Reduce Total Cost of Computing, and More, in a Virtual Environment*, TIPS1099
<http://www.redbooks.ibm.com/abstracts/tips1099.html?Open>
- *Integrated Virtualization Manager for Power Systems Servers*, REDP-4061
<http://www.redbooks.ibm.com/redpieces/abstracts/redp4061.html>
- *IBM Systems Director VMControl Implementation Guide on IBM Power Systems*, SG24-7829
<http://www.redbooks.ibm.com/abstracts/sg247829.html>
- *Server virtualization with IBM PowerVM, The virtualization platform for UNIX, Linux and IBM i clients*
<http://www.ibm.com/systems/power/software/virtualization/index.html>
- *5765-PVS IBM PowerVM Standard Edition and PowerVM Enterprise Edition V2.2* (sales manual)
<http://ibm.co/11ky6J>
- *Power Systems optimization with PowerVM virtualization, Optimize your Power Systems environment with PowerVM virtualization*
<http://www.ibm.com/services/us/en/it-services/power-systems-optimization-with-powervm-virtualization.html>
- *Availability of PowerVM features by Power Systems models*
<http://www.ibm.com/systems/power/software/virtualization/editions/features.html>
- *Hardware Management Console information*
http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/topic/iph8/hwparent_hmc.htm
- IBM Fix Central to check the latest firmware version for HMC
<http://www.ibm.com/support/fixcentral/>
- *IBM Systems Director 6.3.x Information Center*
<http://publib.boulder.ibm.com/infocenter/director/pubs/>

Notices

This information was developed for products and services offered in the U.S.A.

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 give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 states 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. 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. 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 the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

© Copyright International Business Machines Corporation 2013. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document was created or updated on November 26, 2013.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbook@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips1091.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>. The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Active Memory™
AIX®
BladeCenter®
IBM®
POWER®
VMControl™

POWER Hypervisor™
Power Systems™
POWER6®
POWER7®
POWER7 Systems™

PowerLinux™
PowerVM®
Redbooks®
Redbooks (logo)®
Systems Director

The following terms are trademarks of other companies:

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.