

Power Systems

PowerVM Editions

IBM

Power Systems

PowerVM Editions

IBM

Contents

| | |
|---|-----------|
| PowerVM Editions | 1 |
| What's new in PowerVM Editions | 1 |
| PowerVM Editions overview | 1 |
| PowerVM Standard Edition | 5 |
| PowerVM Enterprise Edition | 6 |
| IBM PowerVM, Linux Edition | 7 |
| Entering a PowerVM Editions activation code | 7 |
| Micro-Partitioning technology | 7 |
| Active Memory Sharing | 8 |
| PowerVP Overview | 9 |
| | |
| Notices | 13 |
| Privacy policy considerations | 15 |
| Programming interface information | 15 |
| Trademarks | 15 |
| Terms and conditions | 15 |

PowerVM Editions

Learn about the components and editions of the PowerVM® Editions hardware feature.

What's new in PowerVM Editions

Read about new or changed information in the PowerVM Editions topic collection since the previous update.

October 2016

- The following topics were updated for redundant mover service partitions:
 - “PowerVM Editions overview”
 - “Virtual Network Interface Controller adapter” on page 11
- The following topic is new for the IBM PowerVM, Linux Edition:
 - “IBM PowerVM, Linux Edition” on page 7
- The following topics is updated for the IBM PowerVM, Linux Edition:
 - “PowerVM Editions overview”

October 2015

- The following information is new for virtual Network Interface Controller (vNIC):
 - “Virtual Network Interface Controller adapter” on page 11

October 2014

- Miscellaneous updates were made to this topic collection.

June 2014

- Added information for IBM® Power Systems™ servers that contain the POWER8® processor.

PowerVM Editions overview

Learn about the editions of the PowerVM Editions hardware feature.

The PowerVM Editions hardware feature includes the following editions:

- PowerVM Standard Edition
- PowerVM Enterprise Edition

The following table describes each component of the PowerVM Editions feature and the editions in which each component is included.

Table 1. PowerVM Editions components and editions

| Component | Description | Editions |
|---|--|--|
| Integrated Virtualization Manager (IVM) | You can use the graphical interface of the Virtual I/O Server (VIOS) management partition on some servers that are not managed by a Hardware Management Console (HMC). | <ul style="list-style-type: none">• Standard Edition• Enterprise Edition• IBM PowerVM, Linux Edition |

Table 1. PowerVM Editions components and editions (continued)

| Component | Description | Editions |
|--|--|--|
| Live Partition Mobility | You can migrate an active or inactive AIX®, Linux, or IBM i logical partition from one system to another by using Live Partition Mobility. | <ul style="list-style-type: none"> • Enterprise Edition • IBM PowerVM, Linux Edition |
| Management | Management tools such as HMC, IVM, and PowerVC help to aggregate and manage resources by using a consolidated logical view. | <ul style="list-style-type: none"> • Standard Edition (HMC, Integrated Virtualization Manager, and PowerVC) • Enterprise Edition (HMC, Integrated Virtualization Manager, and PowerVC) • IBM PowerVM, Linux Edition |
| Micro-Partitioning® technology | You can allocate processors to partitions in increments of 0.01, which allows multiple partitions to share the processing power of the system. When the firmware is at level 7.6, or later, micropartitions can be defined as small as 0.05 of a processor and can be changed in increments as small as 0.01 of a processor. A maximum of 20 micropartitions can be created per core. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| N-Port ID Virtualization (NPIV) | NPIV is a standard technology for Fibre Channel networks. With NPIV you can connect multiple partitions to one physical port of a physical Fibre Channel adapter. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| Partition suspend and resume Note: The Suspend/Resume feature for logical partitions is supported on POWER8 processor-based servers when the firmware is at level 8.4.0, or later. | A running AIX, Linux, or IBM i logical partition can be suspended along with its operating system and applications. When a partition is suspended, the state of the partition is saved on persistent storage, and the server resources that were in use by that partition are made available for use by other partitions. At a later time, the operation of the suspended partition and its applications can be resumed. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| PowerVM Active Memory™ Sharing | You can share memory among partitions in a shared memory pool, by using PowerVM Active Memory Sharing. | <ul style="list-style-type: none"> • Enterprise Edition • IBM PowerVM, Linux Edition |
| PowerVP™ | Power Virtualization Performance (PowerVP) is a performance monitoring solution that provides detailed and real-time information about virtualized workloads that are running on Power Systems. 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. | <ul style="list-style-type: none"> • Enterprise Edition • IBM PowerVM, Linux Edition |

Table 1. PowerVM Editions components and editions (continued)

| Component | Description | Editions |
|--------------------------------|--|--|
| PowerVM NovaLink | PowerVM NovaLink is a software interface for virtualization management that you can install on a PowerVM server. PowerVM NovaLink enables highly scalable modern cloud management and deployment for critical enterprise workloads. You can use PowerVM NovaLink to provision large numbers of virtual machines on PowerVM servers quickly and at reduced cost. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| Remote restart | Remote restart is a high availability option for logical partitions. When an error causes a server outage, a partition that is configured for remote restart can be restarted on a different physical server. Sometimes, it might take longer to start the server, in which case remote restart function can be used for faster re-provisioning of the partition. This can be done faster than restarting the server that failed and then restarting the partition. | <ul style="list-style-type: none"> • Enterprise Edition • IBM PowerVM, Linux Edition |
| Shared processor pools | You can assign priorities to partitions, and the hypervisor allocates processing power as needed by the applications. This feature provides automatic nondisruptive balancing of processing power between partitions assigned to shared pools, which results in increased throughput and the potential to reduce processor-based software licensing costs. This feature is supported on POWER6®, POWER7®, and POWER8 processor-based servers. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| Shared storage pools | Shared storage pools provide distributed storage access to all VIOS logical partitions in the cluster. On VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, you can create a cluster of only one VIOS partition that is connected to the same shared storage pool. On VIOS Version 2.2.1.3, or later, you can create a cluster that consists of up to 16 VIOS partitions. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| Single Root I/O Virtualization | Single root I/O virtualization (SR-IOV) is a Peripheral component interconnect express (PCIe) standard architecture that define extensions to PCIe specifications to enable multiple logical partitions running simultaneously within a system to share PCIe devices. The architecture define virtual replicas of PCI functions known as virtual functions (VF). A Partition can connect directly to an SR-IOV adapter VF without going through a virtual intermediary (VI) such as a POWER Hypervisor™ or Virtual I/O Server. This provides for a low latency and lower CPU utilization alternative by avoiding a VI. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |

Table 1. PowerVM Editions components and editions (continued)

| Component | Description | Editions |
|---|--|--|
| Thin provisioning | Thin provisioning helps reduce the amount of unused storage that is allocated to applications or users. Allocation of actual space on the storage subsystem is deferred until the data is written to disk. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| Thick provisioning | Thick provisioning helps reserve storage space for virtual disks, thus ensuring that no failures occur because of lack of storage space. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| VIOS | You can share physical I/O resources on the server between client partitions by using VIOS. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |
| Virtual Network Interface adapter(vNIC) | A logical partition to which single root I/O virtualization (SR-IOV) logical ports are assigned cannot be migrated. To overcome this limitation, a virtual Network Interface Controller (vNIC) can be assigned to a logical partition. A vNIC is a type of virtual adapter that can be configured on client logical partitions to provide a network interface. Each vNIC client adapter is backed by an SR-IOV logical port that is owned by the Virtual I/O Server (VIOS). When the HMC is at version 8.6.0, or later, the firmware is at level 8.6.0, or later, and the VIOS is at version 2.2.5.0, or later, a dedicated vNIC can have multiple SR-IOV logical ports on different physical ports as backing devices, and they can be hosted by the same or different Virtual I/O Servers. | <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition • IBM PowerVM, Linux Edition |

The following table describes the PowerSC™ feature

Table 2. PowerSC feature

| Feature | Description |
|---------|--|
| PowerSC | <p>With HMC Version 7.7.4, or later, and POWER7 processor-based servers with firmware at level 7.4, or later, you can enable Virtual Trusted Platform Module (VTPM) on an AIX logical partition. You can configure partitions on each server by using the HMC, with each partition having its own unique VTPM. The VTPM software implementation checks system boot process, and in association with the AIX Trusted Execution technology provides security and assurance of the boot image on the disk, operating system, and application layers.</p> <p>Trusted Firewall is a feature of PowerSC and you can enable Trusted Firewall in PowerVM. Trusted Firewall is supported on VIOS Version 2.2.1.4, or later. With the Trusted Firewall feature, you can perform inter-VLAN routing functions by using the Security Virtual Machine (SVM) kernel extension. By using Trusted Firewall, partitions that are present on different VLANs of the same server can communicate by using the Shared Ethernet Adapter (SEA).</p> <p>A partition that is enabled for VTPM has the Trusted Boot capability. Trusted Boot is an offering that is supported on the Power® Security and Compliance (PowerSC) Standard Edition. Trusted Boot is a feature of PowerSC. You can enable Trusted Boot in PowerVM. Trusted Boot uses the Virtual Trusted Platform Module (VTPM) software implementation on an AIX logical partition.</p> |

PowerVM Standard Edition

PowerVM Standard Edition includes features designed to allow businesses to increase system utilization; while helping to ensure applications continue to get the resources they need.

The PowerVM Standard Edition includes the following components of the PowerVM Editions hardware feature:

- Micro-Partitioning technology
- N-Port ID Virtualization (NPIV)
- Partition suspend and resume is supported on POWER8 processor-based servers when the firmware is at level 8.4.0, or later.
- PowerVM NovaLink
- Shared processor pools
- Shared storage pools
- Single Root I/O Virtualization (SR-IOV)
- Thin provisioning
- Virtual I/O Server (VIOS)
- Virtual Network Interface Controller adapters

You can use the Standard Edition with IBM Power Systems servers.

In the Hardware Management Console (HMC) environment, the IVM is disabled because the system is managed by the HMC. For servers that are not managed by an HMC, the VIOS becomes the management partition and provides the IVM to help you manage the system.

An activation code is required to enable the PowerVM Editions hardware feature. When you specify the feature with the initial system order, the firmware is activated to support the components of the feature. If you order the feature separately from the server, you can enter the code by using the HMC, IVM.

Related information:

Entering the activation code for PowerVM Editions with the Integrated Virtualization Manager

Entering the activation code for PowerVM Editions using the HMC version 7

PowerVM Enterprise Edition

PowerVM Enterprise Edition includes all the features of PowerVM Standard Edition and two more capabilities called Active Memory Sharing and Live Partition Mobility.

PowerVM Enterprise Edition includes the following components of the PowerVM Editions hardware feature:

- PowerVM Active Memory Sharing
- Micro-Partitioning technology
- Live Partition Mobility

Note: On POWER7 processor-based servers with firmware at level 7.3, or later, you can evaluate Live Partition Mobility at no charge by using Trial Live Partition Mobility. Trial Live Partition Mobility requires an activation code. You must have a valid license and an active Software Maintenance contract for PowerVM Standard Edition to use Trial Live Partition Mobility.

- Partition suspend and resume is supported on POWER8 processor-based servers when the firmware is at level 8.4.0, or later.
- N-Port ID Virtualization (NPIV)
- PowerVM NovaLink
- Remote restart
- Shared processor pools
- Shared storage pools
- Single Root I/O Virtualization (SR-IOV)
- Thin provisioning
- Virtual I/O Server (VIOS)
- Power Virtualization Performance (PowerVP)

Note: PowerVP requires firmware version 7.7, or later.

- Virtual Network Interface Controller adapters

You can use the Enterprise Edition with IBM Power Systems servers.. Like the Standard Edition, in a Hardware Management Console (HMC) environment, the IVM is disabled because the system is managed by the HMC. For servers that are not managed by an HMC, the VIOS becomes the management partition and provides the IVM to help you manage the system.

An activation code is required to enable the PowerVM Editions hardware feature. When you specify the feature with the initial system order, the firmware is activated to support the components of the feature. If you order the feature separately from the server, you can enter the code by using the HMC, or the IVM. This code differs from the PowerVM Standard Edition activation code. For example, you might have previously purchased the Standard Edition and enabled the feature on the system. Now you want the ability to move partitions from one system to another. To do so, you must purchase the Enterprise Edition and enter the activation code for the Enterprise Edition.

Related information:

Entering the activation code for PowerVM Editions with the Integrated Virtualization Manager

Entering the activation code for PowerVM Editions using the HMC version 7

IBM PowerVM, Linux Edition

IBM PowerVM, Linux Edition includes all the features of PowerVM Enterprise Edition, and runs only on Linux based servers.

IBM PowerVM, Linux Edition includes the following components of the PowerVM Editions hardware feature:

- PowerVM Active Memory Sharing
- Micro-Partitioning technology
- Live Partition Mobility

Note: On POWER7 processor-based servers with firmware at level 7.3, or later, you can evaluate Live Partition Mobility at no charge by using Trial Live Partition Mobility. Trial Live Partition Mobility requires an activation code. You must have a valid license and an active Software Maintenance contract for PowerVM Standard Edition to use Trial Live Partition Mobility.

- Partition suspend and resume is supported on POWER8 processor-based servers when the firmware is at level 8.4.0, or later.
- N-Port ID Virtualization (NPIV)
- PowerVM NovaLink
- Remote restart
- Shared processor pools
- Shared storage pools
- Single Root I/O Virtualization (SR-IOV)
- Thin provisioning
- Virtual I/O Server (VIOS)
- Power Virtualization Performance (PowerVP)

Note: PowerVP requires firmware version 7.7, or later.

- Virtual Network Interface Controller adapters

Entering a PowerVM Editions activation code

You can enter the activation code for PowerVM Editions by using the Integrated Virtualization Manager, or the Hardware Management Console (HMC). A Trial Live Partition Mobility code can be entered only by using the HMC.

Micro-Partitioning technology

When you enable the Micro-Partitioning technology, you can configure multiple partitions to share system processing power.

All processors that are not dedicated to specific partitions are placed in the shared processor pool that is managed by the hypervisor. Partitions that are set to use shared processors can use the shared processor pool. When the firmware is at level 7.5, or earlier, you can set a partition that uses shared processors to use as little as 0.10 processing unit, which is one-tenth of the processing capacity of a single processor. When the firmware is at level 7.6, or later, you can set a partition that uses shared processors to use as little as 0.05 processing units, which is approximately a 20th of the processing capacity of a single processor. You can specify the number of processing units to be used by a shared processor partition down to a hundredth of a processing unit. This ability to assign fractions of processing units to partitions and allowing partitions to share processing units is called Micro-Partitioning technology.

The Micro-Partitioning technology allows for increased overall use of system resources by automatically applying only the required amount of processor resource needed by each partition. The hypervisor can

automatically and continually adjust the amount of processing capacity allocated to each logical partition (with shared processors) based on system demand. You can set a shared processor partition so that, if the partition requires more processing capacity than its assigned number of processing units, the partition can use unused processing units from the shared processor pool.

The Micro-Partitioning technology is part of the PowerVM Editions hardware feature and is available on all IBM Power Systems servers.

The Micro-Partitioning technology is supported by all the operating environments available on POWER5 , POWER6 , POWER7 , and POWER8 processor-based servers, including:

- AIX
- IBM i
- Linux
- Virtual I/O Server

Active Memory Sharing

When you enable the PowerVM Active Memory Sharing technology, you can configure multiple partitions to share system memory.

Note: For the initial introduction of the POWER8 8286-41A, 8286-42A, 8286-42A, 8247-21L, and 8247-22L Power Systems servers, Live Partition Mobility is not supported when used with Active Memory Sharing configurations.

Partitions can share the memory in the shared memory pool by using the PowerVM Active Memory Sharing technology (or shared memory). Instead of assigning a dedicated amount of physical memory to each partition that uses shared memory (hereafter referred to as *shared memory partitions*), the hypervisor constantly provides the physical memory from the shared memory pool to the shared memory partitions as needed.

The hypervisor provides portions of the shared memory pool that are not currently being used by shared memory partitions to other shared memory partitions that need to use the memory. When a shared memory partition needs more memory than the current amount of unused memory in the shared memory pool, the hypervisor stores a portion of the memory that belongs to the shared memory partition in auxiliary storage. Access to the auxiliary storage is provided by a Virtual I/O Server (VIOS) partition. When the operating system attempts to access data that is in auxiliary storage, the hypervisor directs a VIOS to retrieve the data from the auxiliary storage and write it to the shared memory pool so that the operating system can access the data.

The PowerVM Active Memory Sharing technology is available with the product, which also includes the license for the VIOS software.

Related information:

Partitioning

Virtual I/O Server

Virtual I/O Server and Integrated Virtualization Manager commands

Thin provisioning

In Virtual I/O Server Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later, a client virtual-SCSI device can be thin-provisioned for better storage space utilization. In a thin-provisioned device, the used storage space might be greater than the actual used storage space. If the blocks of storage space in a thin-provisioned device are unused, the device is not entirely backed up by physical storage space. With thin-provisioning, the storage capacity of the storage pool can be exceeded.

Related information:

Thick provisioning

In VIOS Version 2.2.1.3, or later, you can use a thick-provisioned virtual disk. In a thick-provisioned virtual disk, you can allocate or reserve storage space while initially provisioning the virtual disk. The allocated storage space for the thick-provisioned virtual disk is guaranteed. This operation ensures that there are no failures because of lack of storage space. By using thick-provisioning, initial storage in the virtual disk is faster because disk blocks might be allocated to the disk during initial provisioning of the virtual disk.

PowerVP Overview

Power Virtualization Performance (PowerVP) is a performance monitoring solution that provides detailed and real-time information about virtualized workloads that are running on Power Systems. 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 that some action or resource adjustment should be considered. The colors and thresholds are customizable as the thresholds are installation-dependent. Metrics include nodes, processor modules, cores, Powerbus 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. It provides a recording feature for storing performance information with DVR-like functions such as play, fast forward, rewind, jump, pause, or stop. You can find performance bottlenecks by playing back the recorded data at any point in time.

Single root I/O virtualization

Single root I/O virtualization (SR-IOV) is a Peripheral component interconnect express (PCIe) standard architecture that define extensions to PCIe specifications to enable multiple partitions running simultaneously within a system to share PCIe devices. The architecture define virtual replicas of PCI functions known as virtual functions (VF). A Partition can connect directly to an SR-IOV adapter VF without going through a virtual intermediary (VI) such as a POWER Hypervisor or Virtual I/O Server. This provides for a low latency and lower CPU utilization alternative by avoiding a VI.

SR-IOV capable adapter might be assigned to a partition in a dedicated or shared mode. The management console provides an interface to enable SR-IOV adapter in a shared mode. An SR-IOV capable adapter in a shared mode is assigned to the POWER Hypervisor for management of the adapter and provisioning of adapter resources to partitions. The management console, in conjunction with the POWER Hypervisor, provides the ability to manage the adapter's physical Ethernet ports and logical ports. To connect a partition to an SR-IOV Ethernet adapter VF, create an SR-IOV Ethernet logical port for the partition. When you create an Ethernet logical port for a partition, select the adapter physical Ethernet

port to connect to the partition and specify the resource requirements for the logical port. Each partition can have one or more logical ports from each SR-IOV adapter in shared mode. The number of logical ports for all configured partitions cannot exceed the adapter logical port limit.

To create an SR-IOV Ethernet logical port for a partition, use one of the following methods:

- Create an Ethernet logical port when you create a partition.
- Add an Ethernet logical port to a partition profile, shutdown the partition, and reactivate the partition by using the partition profile.
- Add an Ethernet logical port to a running partition by using dynamic partitioning.

Note: An SR-IOV adapter does not support Live Partition Mobility unless the VF is assigned to a shared Ethernet adapter.

When you activate a partition, the logical ports in the partition profile are considered to be a required resource. If the physical adapter resources required by the logical port are not available, the partition cannot be activated. However, logical ports can be removed dynamically from other partition to make the required resources available to the partition that is activated.

For an SR-IOV adapter in shared mode, the physical port switch mode can be configured in Virtual Ethernet Bridge (VEB) or Virtual Ethernet Port Aggregator (VEPA) mode. If the switch mode is configured in VEB mode, the traffic between the logical ports are not visible to the external switch. If the switch mode is configured in VEPA mode, the traffic between logical ports must be routed back to the physical port by the external switch. Before enabling the physical port switch in VEPA mode, ensure that the switch attached to the physical port is supported and is enabled for reflective relay.

When you create an Ethernet logical port, you can select a promiscuous permission to allow the logical port to be configured as a promiscuous logical port by the partition. A promiscuous logical port receives all unicast traffic with a destination address that does not match the address of one of the other logical ports configured for the same physical port. The number of logical ports configured for partitions, active or shutdown, on a physical port will be limited to minimize potential performance impact due to overhead associated with promiscuous logical ports. The management console indicates the number of logical ports on the physical port that are allowed to have a promiscuous permission setting.

When bridging between virtual Ethernet adapters and a physical Ethernet adapter, an SR-IOV Ethernet logical port might be used as the physical Ethernet adapter to access the outside network. When a logical port is configured as the physical Ethernet adapter for bridging, the logical port must have the promiscuous permission enabled. For example, if you create a logical port for a Virtual I/O Server logical partition and to use the logical port as the physical adapter for the shared Ethernet adapter, you must select the promiscuous permission for the logical port.

Configuration requirements

Consider the following configuration requirements when an Ethernet logical port is used as the physical Ethernet device for shared Ethernet adapter bridging:

- When there is a requirement to divert all network traffic to flow through an external switch, consider the following requirements:
 - The POWER Hypervisor virtual switch must be set to the VEPA switching mode and the SR-IOV Ethernet adapter physical port switch mode must also be set to the VEPA switching mode.
 - In addition, the logical port is the only logical port configured for the physical port.
- When you create an Ethernet logical port you can also specify a capacity value. The capacity value specifies the desired capacity of the logical port as a percentage of the capability of the physical port. The capacity level value determines the amount of resources that are assigned to the logical port from the physical port. The assigned resources determine the minimum capability of the logical port. The physical port resources that are not used by other logical ports might be temporarily used by the

logical port when the logical port exceeds its assigned resources to allow additional capability. System or network limitations can influence the amount of throughput a logical port can actually achieve. The maximum capacity that can be assigned to a logical port is 100%. The sum of the capacity values for all the configured logical ports on a physical port must be less than or equal to 100%. To minimize the configuration effort while adding additional logical ports, you can reserve some of the physical port capacity for additional logical ports.

- When an Ethernet logical port is used as a physical adapter for bridging virtual Ethernet adapters, the parameter values such as the number of client virtual adapters and expected throughput must be considered when choosing a capacity value.
- The Ethernet logical ports allow the logical port to run diagnostics on the adapter and physical port. Select this permission only while running the diagnostics by using the logical port.

Remote restart

Remote restart is a high availability option for logical partitions. When an error causes a server outage, a partition that is configured with the remote restart capability can be restarted on a different physical server. Sometimes, it might take longer to start the server, in which case the remote restart feature can be used for faster re-provisioning of the partition. This operation completes faster as compared to restarting the server that failed and then restarting the partition.

Remote restart is supported on POWER7 and POWER8 processor-based systems. The simplified version of the remote restart feature is supported when the firmware is at level 8.2.0, or later, and the HMC is at version 8.2.0, or later. It is recommended to use the simplified version of the remote restart feature when the firmware is at level 8.2.0, or later, and the HMC is at version 8.2.0, or later.

The following are the characteristics of the remote restart feature:

- The remote restart feature is not a Suspend/Resume or migration operation of the partition that preserves the active running state of the partition. During the remote restart operation, the logical partition is shut down and then restarted on a different system.
- The remote restart feature preserves the resource configuration of the partition. If processors, memory or I/O are added or removed while the partition is running, the remote restart operation activates the partition with the most recent configuration.

The remote restart feature requires a reserved storage device that is assigned to each partition. You must manage a reserved storage device pool on both the source and the destination servers, and maintain a record of the device that is assigned to each partition. The simplified remote restart feature does not require a reserved storage device that is assigned to each partition.

The remote restart feature (including the simplified version) is not supported from the HMC for logical partitions that are co-managed by the HMC and PowerVM NovaLink. However, you can run simplified remote restart operations by using PowerVC with PowerVM NovaLink.

Related information:

Configuration requirements and restrictions for the remote restart capability of a logical partition

Remote restart states

Enabling or disabling the remote restart capability of a logical partition by using the HMC

Remote restart of a logical partition

Virtual Network Interface Controller adapter

A logical partition to which single root I/O virtualization (SR-IOV) logical ports that are assigned cannot be migrated. To overcome this limitation, a virtual Network Interface Controller (vNIC) can be assigned to a logical partition. A vNIC is a type of virtual adapter that can be configured on client logical partitions to provide a network interface. Each vNIC client adapter is backed by an SR-IOV logical port that is owned by the Virtual I/O Server (VIOS).

When the VIOS is at version 2.2.5.0, and the firmware is at level 8.6.0, or later, you can specify redundant mover service partitions for a single partition mobility operation. Redundancy of mover service partitions provides for better reliability of partition mobility operations in the event of VIOS failure, some HMC failures, or network failures.

Virtual NICs improve the performance of virtual networks. Dedicated vNICs enables the vNIC assigned to the partition to leverage the Quality of Service (QoS) capability that is available with SR-IOV adapters. A logical partition to which vNIC adapters are assigned can be migrated when the destination server also supports vNIC adapters.

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.

Licensees 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 Setting up the virtualization environment publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM AIX Version 7.2, IBM AIX Version 7.1, IBM AIX Version 6.1, IBM i 7.3, and IBM Virtual I/O Server Version 2.2.4.20.

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](http://www.ibm.com/legal/copytrade.shtml).

Linux is a registered trademark of Linus Torvalds 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