# PowerVM Editions

## Operations Guide

ESCALA

REFERENCE
86 A1 40EV 03

# ESCALA

# PowerVM Editions

## Operations Guide

Hardware

July 2008

BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

REFERENCE
86 A1 40EV 03

Printed in France

## Trademarks and Acknowledgements

*The information in this document is subject to change without notice. Bull will not be liable for errors contained herein, or for incidental or consequential damages in connection with the use of this material.*

# Contents

# About this publication

This publication provides system operators and administrators with information about installing, configuring, managing, and using the PowerVM™ Editions (or Advanced POWER Virtualization) feature. PowerVM Editions includes the Micro-Partitioning™ technology, Virtual I/O Server, Integrated Virtualization Manager, Live Partition Mobility, Partition Load Manager for AIX®, and Lx86.

For information about the accessibility features of this product, for users who have a physical disability, see “Accessibility features,” on page 299.

## How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information. If you have any comments about this publication, send your comments to us. Be sure to include the name of the book and the specific location of the text you are commenting on (for example, a page number or table number).

# Chapter 1. PowerVM Editions

Learn about the components and editions of the PowerVM Editions (formerly known as Advanced POWER™ Virtualization) hardware feature.

The PowerVM Editions hardware feature includes the following components to enhance the virtualization capabilities of your system:

- Micro-Partitioning technology
- Virtual I/O Server
- Integrated Virtualization Manager
- Live Partition Mobility
- Partition Load Manager for AIX
- Lx86

The PowerVM Editions hardware feature includes the following editions:

- PowerVM Express Edition
- PowerVM Standard Edition
- PowerVM Enterprise Edition

The following table describes each component of the PowerVM Editions feature, the editions in which each component is included, and the processor-based hardware on which each component is available.

*Table 1. PowerVM Editions components, editions, and hardware support*

| Component | Description | Editions | Hardware |
|---|---|---|---|
| Micro-Partitioning technology | The ability to allocate processors to logical partitions in increments of 0.1 allowing multiple logical partitions to share the system's processing power. | • Express Edition<br>• Standard Edition<br>• Enterprise Edition | • POWER6™<br>• POWER5™ |
| Virtual I/O Server | Software that facilitates the sharing of physical I/O resources between client logical partitions within the server. | • Express Edition<br>• Standard Edition<br>• Enterprise Edition | • POWER6<br>• POWER5 |
| Integrated Virtualization Manager | The graphical interface of the Virtual I/O Server management partition on some servers that are not managed by an Hardware Management Console. | • Express Edition<br>• Standard Edition<br>• Enterprise Edition | • POWER6<br>• POWER5 |
| Live Partition Mobility | The ability to migrate an active or inactive AIX or Linux® logical partition from one system to another. | Enterprise Edition | POWER6 |
| Partition Load Manager | Software that provides processor and memory resource management and monitoring across AIX logical partitions within a single central processor complex. | Standard Edition | POWER5 |
| Lx86 | A product that makes a POWER system compatible with x86 applications. This extends the application support for Linux on POWER systems, allowing applications that are available on x86 but not on POWER systems to be run on the POWER system. | • Express Edition<br>• Standard Edition<br>• Enterprise Edition | POWER6 running SUSE or Red Hat Linux |

## PowerVM Editions

Learn about the PowerVM Express Edition, the PowerVM Standard Edition, and the PowerVM Enterprise Edition.

### PowerVM Express Edition

The PowerVM Express Edition includes the following components of the PowerVM Editions (formerly known as Advanced POWER Virtualization) hardware feature:

- Micro-Partitioning technology
- Virtual I/O Server
- Integrated Virtualization Manager
- Lx86

You can use the Express Edition in restricted environments at no cost. The Express Edition is available on some POWER6 processor-based servers. To use the Express Edition, the system must be managed by the Integrated Virtualization Manager.

With the Express Edition, you can create up to two client logical partitions that use virtual Small Computer System Interface (SCSI) and shared processors, also called Micro-Partitioning technology. If you want to create more than two client logical partitions that use shared processors or virtual SCSI, then you must purchase either the Standard or Enterprise Edition and enter the activation code. (Virtual Ethernet is available to all client logical partitions.)

### PowerVM Standard Edition

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

- Micro-Partitioning technology
- Virtual I/O Server
- Integrated Virtualization Manager
- Partition Load Manager for AIX
- Lx86

You can use the Standard Edition with POWER5 and POWER6 processor-based servers. However, Partition Load Manager is available only on POWER5-processor based systems.

In HMC environments, the Integrated Virtualization Manager is disabled because the system is managed by the HMC. For servers that are not managed by an HMC, the Virtual I/O Server becomes the management partition and provides the Integrated Virtualization Manager to help you manage the system.

For some servers, the Standard Edition requires an activation code to enable the PowerVM Editions hardware feature. When you specify the PowerVM Editions hardware 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 using the HMC or the Integrated Virtualization Manager. For BladeCenter® blade servers, the server includes the Standard Edition; you do not need an activation code.

### PowerVM Enterprise Edition

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

- Micro-Partitioning technology
- Virtual I/O Server
- Integrated Virtualization Manager
- Live Partition Mobility
- Lx86

You can use the Enterprise Edition with POWER6 processor-based servers.Like the Standard Edition, in HMC environments, the Integrated Virtualization Manager is disabled because the system is managed by the HMC. For servers that are not managed by an HMC, the Virtual I/O Server becomes the management partition and provides the Integrated Virtualization Manager to help you manage the system.

The Enterprise Edition requires an activation code 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 using the HMC or the Integrated Virtualization Manager. This is a different code from the Standard Edition. For example, you might have previously purchased the Standard Edition and enabled the feature on the system. Now you want the ability to migrate logical partitions from one system to another. To do this, you must purchase the Enterprise Edition and enter the activation code for the Enterprise Edition.

## Micro-Partitioning technology

Learn about Micro-Partitioning technology, its availability, and supported hardware.

The PowerVM Editions (or Advanced POWER Virtualization) feature includes firmware enablement for Micro-Partitioning technology. Micro-Partitioning technology is the ability to allocate processors to logical partitions in increments of 0.01, with a minimum of 0.1, allowing multiple logical partitions to share the system's processing power.

Micro-Partitioning technology is included in the PowerVM Express Edition, the PowerVM Standard Edition, and the PowerVM Enterprise Edition.

Micro-Partitioning technology is available for all POWER5 and POWER6 processor-based servers.

## Virtual I/O Server

Learn about the Virtual I/O Server, its availability, and supported hardware.

The PowerVM Editions (or Advanced POWER Virtualization) hardware feature includes the installation media for the Virtual I/O Server software. The Virtual I/O Server facilitates the sharing of physical I/O resources client logical partitions within the server. The Virtual I/O Server contains one charge unit per activated processor, including software maintenance.

The Virtual I/O Server is included in the PowerVM Express Edition in restricted environments, in the PowerVM Standard Edition, and in the PowerVM Enterprise Edition.

Virtual I/O Server is available for all POWER5 and POWER6 processor-based systems.

## Integrated Virtualization Manager

Learn about the Integrated Virtualization Manager, its availability, and supported hardware.

The PowerVM Editions (or Advanced POWER Virtualization) feature includes the installation media for the Virtual I/O Server software. The Virtual I/O Server facilitates the sharing of physical I/O resources between client logical partitions within the server.

For some servers that are not managed by a Hardware Management Console (HMC), the Virtual I/O Server becomes the management partition and provides a systems management interface, called the Integrated Virtualization Manager, to help you manage the system.

The Virtual I/O Server (including the Integrated Virtualization Manager) is a licensed software component of the PowerVM Editions feature. It contains one charge unit per activated processor, including software maintenance.

The Integrated Virtualization Manager is included in the PowerVM Express Edition, the PowerVM Standard Edition, and the PowerVM Enterprise Edition.

The Integrated Virtualization Manager is available for the following server models:
- 7/10
- 7/20
- 04E/8A
- 03E/4A
- 04E/8A
- 105/1A
- 105/10
- 115/20
- 135/50
- 155/05
- 165/61
- 315/2A
- 335/5A
- JS/21
- JS/22
- JS/12

## Live Partition Mobility

Learn about Live Partition Mobility, its availability, and supported hardware.

Live Partition Mobility provides the ability to migrate an active or inactive AIX or Linux logical partition from one system to another. Active Partition Mobility refers to moving a running logical partition, including its operating system and applications, from one system to another. The logical partition and the applications running on the logical partition do not need to be shut down. Inactive Partition Mobility refers to moving a powered off logical partition from one system to another.

For systems that are managed by an HMC, you can migrate a powered off or powered on partition to a different system that is managed by the same HMC. For systems that are managed by the Integrated Virtualization Manager, you can migrate a powered off or powered on partition to a different system that is managed by a different Integrated Virtualization Manager.

Live Partition Mobility is included in the PowerVM Enterprise Edition.

Live Partition Mobility is available on some POWER6 processor-based servers, the JS/22 and JS/12 Express.

## Partition Load Manager for AIX

Learn about Partition Load Manager, its availability, and supported hardware.

Partition Load Manager provides processor and memory resource management and monitoring across AIX logical partitions within a single central processor complex.

Partition Load Manager for AIX is a licensed software component of the PowerVM Editions (or Advanced POWER Virtualization) feature. It contains one charge unit per activated processor, including software maintenance.

Partition Load Manager for AIX is included in the PowerVM Standard Edition.

Partition Load Manager is not available on POWER6 processor-based servers.

Partition Load Manager is available on some POWER5 processor-based servers

## PowerVM Lx86

Learn about PowerVM Lx86, its availability, and supported hardware.

The PowerVM Editions (formerly known as Advanced POWER Virtualization) hardware feature includes PowerVM Lx86. Lx86 is a dynamic, binary translator that allows Linux applications (compiled for Linux on Intel®) to run without change, alongside local Linux on POWER applications. Lx86 makes this possible by dynamically translating x86 instructions to POWER and caching them to enhance translation performance. In addition, Lx86 maps Linux on Intel system calls to Linux on Power system calls. No modifications or recompilations of the x86 Linux applications are needed.

Lx86 creates a virtual x86 environment, within which, the Linux on Intel applications can run. Currently, a virtual Lx86 environment supports SUSE or Red Hat Linux x86 distributions. The translator and the virtual environment run strictly within the user-space. No modifications to the POWER kernel are required. Lx86 does not run the x86 kernel on the POWER machine. The Lx86 virtual environment is not a virtual machine. Instead, x86 applications are encapsulated so the operating environment appears to be Linux on x86, even though the underlying system is a Linux on POWER system.

Lx86 is included in the PowerVM Express Edition, PowerVM Standard Edition, and in the PowerVM Enterprise Edition.

For more information about Lx86, see PowerVM Lx86 for x86 Linux Applications Administration Guide.

# Chapter 2. The Micro-Partitioning technology

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

All processors that are not dedicated to specific logical partitions are placed in the shared processor pool that is managed by the hypervisor. Logical partitions that are set to use shared processors can use the shared processor pool. You can set a logical partition that uses shared processors to use as little as 0.10 processing units, which is approximately one-tenth of the processing capacity of a single processor. You can specify the number of processing units to be used by a shared processor logical partition down to the hundredth of a processing unit. This ability to assign fractions of processing units to logical partitions and allowing logical partitions to share processing units is called the 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 logical 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 logical partition so that, if the logical partition requires more processing capacity than its assigned number of processing units, the logical 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 for all POWER5 and POWER6 processor-based servers.

The Micro-Partitioning technology is supported by the following operating environments:
- AIX 5.3 + APAR IY58321 or later
- Linux
- Virtual I/O Server version 1.0 or later (version 1.2 or later for the Integrated Virtualization Manager)

## Shared processors

*Shared processors* are physical processors whose processing capacity is shared among multiple logical partitions. The ability to divide physical processors and share them among multiple logical partitions is known as the *Micro-Partitioning* technology.

**Note:** For some models, the Micro-Partitioning technology is an option for which you must obtain and enter a PowerVM Editions (or Advanced POWER Virtualization) activation code.

By default, all physical processors that are not dedicated to specific logical partitions are grouped together in a *shared processor pool*. You can assign a specific amount of the processing capacity in this shared processor pool to each logical partition that uses shared processors. Some models allow you to use the HMC to configure multiple shared processor pools. These models have a *default shared processor pool* that contains all the processors that do not belong to logical partitions that use dedicated processors or logical partitions that use other shared processor pools. The other shared processor pools on these models can be configured with a maximum processing unit value and a reserved processing unit value. The maximum processing unit value limits the total number of processing unit that can be used by the logical partitions in the shared processor pool. The reserved processing unit value is the number of processing units that are reserved for the use of uncapped logical partitions within the shared processor pool.

You can assign partial processors to a logical partition that uses shared processors. A minimum of 0.10 processing units can be configured for any partition that uses shared processors. Processing units are a unit of measure for shared processing power across one or more virtual processors. One shared processing unit on one virtual processor accomplishes approximately the same work as one dedicated processor.

Some server models allow logical partitions to use only a portion of the total active processors on the managed system, so you are not always able to assign the full processing capacity of the managed system to logical partitions. This is particularly true for server models with one or two processors, where a large portion of processor resources is used as overhead.

Shared processors are assigned to logical partitions using partition profiles. For more information about how partition profiles are used to specify resource configurations, see Partition profile.

Partitions that use shared processors can have a sharing mode of capped or uncapped. An *uncapped logical partition* is a logical partition that can use more processor power than its assigned processing capacity. The amount of processing capacity that an uncapped logical partition can use is limited only by the number of virtual processors assigned to the logical partition or the maximum processing unit allowed by the shared processor pool that the logical partition uses. In contrast, a *capped logical partition* is a logical partition that cannot use more processor power than its assigned processing units.

For example, logical partitions 2 and 3 are uncapped logical partitions, and logical partition 4 is a capped logical partition. Logical partitions 2 and 3 are each assigned 3.00 processing units and four virtual processors. Logical partition 2 currently uses only 1.00 of its 3.00 processing units, but logical partition 3 currently has a workload demand that requires 4.00 processing units. Because logical partition 3 is uncapped and has four virtual processors, the server firmware automatically allows logical partition 3 to use 1.00 processing units from logical partition 2. This increases the processing power for logical partition 3 to 4.00 processing units. Soon afterwards, logical partition 2 increases its workload demand to 3.00 processing units. The server firmware therefore automatically returns 1.00 processing units to logical partition 2 so that logical partition 2 can use its full, assigned processing capacity once more. Logical partition 4 is assigned 2.00 processing units and three virtual processors, but currently has a workload demand that requires 3.00 processing units. Because logical partition 4 is capped, logical partition 4 cannot use any unused processing units from logical partitions 2 or 3. However, if the workload demand of logical partition 4 decreases below 2.00 processing units, logical partitions 2 and 3 could use any unused processing units from logical partition 4.

By default, logical partitions that use shared processors are capped logical partitions. You can set a logical partition to be an uncapped logical partition if you want the logical partition to use more processing power than its assigned amount.

Although an uncapped logical partition can use more processor power than its assigned processing capacity, the uncapped logical partition can never use more processing units than its assigned number of virtual processors. Also, the logical partitions that use a shared processor pool can never use more processing units than the maximum processing units configured for the shared processor pool.

If multiple uncapped logical partitions need additional processor capacity at the same time, the server can distribute the unused processing capacity to all uncapped logical partitions. This distribution process is determined by the uncapped weight of each of the logical partitions.

*Uncapped weight* is a number in the range of 0 through 255 that you set for each uncapped partition in the shared processor pool. On the HMC, you can choose from any of the 256 possible uncapped weight values. The Integrated Virtualization Manager limits you to only one of several different uncapped weight values. By setting the uncapped weight (255 being the highest weight), any available unused capacity is distributed to contending logical partitions in proportion to the established value of the uncapped weight. The default uncapped weight value is 128.

For example, logical partition 2 has an uncapped weight of 100, and logical partition 3 has an uncapped weight of 200. If logical partitions 2 and 3 both require additional processing capacity, logical partition 3 would receive two additional processing units for every additional processing unit that logical partition 2 receives.

**Related information**

Logical Partitioning Guide

## Virtual processors

A *virtual processor* is a representation of a physical processor core to the operating system of a logical partition that uses shared processors.

When you install and run an operating system on a server that is not partitioned, the operating system calculates the number of operations that it can perform concurrently by counting the number of processors on the server. For example, if you install an operating system on a server that has eight processors, and each processor can perform two operations at a time, the operating system can perform 16 operations at a time. In the same way, when you install and run an operating system on a logical partition that uses dedicated processors, the operating system calculates the number of operations that it can perform concurrently by counting the number of dedicated processors that are assigned to the logical partition. In both cases, the operating system can easily calculate how many operations it can perform at a time by counting the whole number of processors that are available to it.

However, when you install and run an operating system on a logical partition that uses shared processors, the operating system cannot calculate a whole number of operations from the fractional number of processing units that are assigned to the logical partition. The server firmware must therefore represent the processing power available to the operating system as a whole number of processors. This allows the operating system to calculate the number of concurrent operations that it can perform. A *virtual processor* is a representation of a physical processor to the operating system of a logical partition that uses shared processors.

The server firmware distributes processing units evenly among the virtual processors assigned to a logical partition. For example, if a logical partition has 1.80 processing units and two virtual processors, each virtual processor has 0.90 processing units supporting its workload.

There are limits to the number of processing units that you can have for each virtual processor. The minimum number of processing units that you can have for each virtual processor is 0.10 (or ten virtual processors for every processing unit). The maximum number of processing units that you can have for each virtual processor is always 1.00. This means that a logical partition cannot use more processing units than the number of virtual processors that it is assigned, even if the logical partition is uncapped.

A logical partition generally performs best if the number of virtual processors is close to the number of processing units available to the logical partition. This lets the operating system manage the workload on the logical partition effectively. In certain situations, you might be able to increase system performance slightly by increasing the number of virtual processors. If you increase the number of virtual processors, you increase the number of operations that can run concurrently. However, if you increase the number of virtual processors without increasing the number of processing units, the speed at which each operation runs will decrease. The operating system also cannot shift processing power between processes if the processing power is split between many virtual processors.

Virtual processors are assigned to logical partitions using partition profiles. For more information on how partition profiles are used to specify resource configurations, see Partition profile.

**Related information**

Logical Partitioning Guide
This publication contains information about logical partition profiles.

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

You can enter the activation code for PowerVM Editions (or Advanced POWER Virtualization) using the Integrated Virtualization Manager.

The code level for the Integrated Virtualization Manager must be at version 1.5, or later, to perform the following procedure. For instructions about how to view and update the current code level, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.

Whether you need to enter an activation code depends on your edition of the PowerVM Editions feature and the hardware on which you plan to enable it. The following table summarizes the requirements.

*Table 2. Activation code requirements*

| | systems | BladeCenter blade servers |
|---|---|---|
| **PowerVM Express Edition** | No activation code is required. | The Express Edition is not available on blade servers. |
| **PowerVM Standard Edition** | The PowerVM Editions activation code is required. | No activation code is required. The Standard Edition is included with the blade server. |
| **PowerVM Enterprise Edition** | The PowerVM Editions activation code is required.<br>**Note:** If you already have the Standard Edition enabled, you must enter a separate, additional activation code for the Enterprise Edition. | The PowerVM Editions activation code is required. |

For detailed information about the PowerVM Editions editions, see "PowerVM Editions" on page 2.

Before you start, verify that you have access to the Integrated Virtualization Manager. For instructions, see "Installing the Virtual I/O Server and enabling the Integrated Virtualization Manager on systems servers" on page 169.

To enter the activation code in the Integrated Virtualization Manager, complete the following tasks:

1. From the **IVM Management** menu, click **Enter PowerVM Editions Key**. The Enter PowerVM Editions Key window is displayed.
2. Enter your activation code for PowerVM Editions and click **Apply**.

You can now create more than two client logical partitions that use virtual I/O or shared processors.

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

Use these instructions to enter the PowerVM Editions (or Advanced POWER Virtualization) activation code using the Hardware Management Console (HMC) version 7, or later.

If PowerVM Editions is not enabled on your system, you can use the HMC to enter the activation code that you received when you ordered the feature.

Use the following procedure to enter the activation code for the PowerVM Standard Edition and the PowerVM Enterprise Edition. For information about the PowerVM Editions, see "PowerVM Editions" on page 2.

To enter your activation code, follow these steps:

1. In the Navigation area, expand **Systems Management**.
2. Select **Servers**.
3. In the contents area, select the managed system on which you plan to use PowerVM Editions. For example, this might be the system on which you plan to install the Virtual I/O Server, or it might be the system in which you plan to use the Micro-Partitioning technology.
4. Click **Tasks** and select **Capacity on Demand (CoD)** → **Advanced POWER Virtualization** → **Enter Activation Code**.
5. Enter your activation code and click **OK**.

## Entering the activation code for PowerVM Editions using the HMC version 6

The PowerVM Editions (or Advanced POWER Virtualization) activation code is required to install and configure the Virtual I/O Server. You can enter the code using the Hardware Management Console (HMC).

If the PowerVM Editions feature is not enabled on your system, you must use the HMC to enter the activation code that you received when you ordered the feature. This activation code also enables the Micro-Partitioning technology on the system.

To enter your activation code, follow these steps:

1. From the HMC, select the managed system.
2. Select **Manage On Demand Activations**.
3. Select **Virtualization Engine Technologies**.
4. Select **Enter Activation Code**. Type your activation code.

# Chapter 3. Virtual I/O Server

Manage the Virtual I/O Server and client logical partitions using the Hardware Management Console (HMC) and the Virtual I/O Server command-line interface.

The PowerVM Editions (or Advanced POWER Virtualization) feature includes the installation media for the Virtual I/O Server software. The Virtual I/O Server facilitates the sharing of physical I/O resources between client logical partitions within the server.

When you install the Virtual I/O Server in a logical partition on a system that is managed by the HMC, you can use the HMC and the Virtual I/O Server command-line interface to manage the Virtual I/O Server and client logical partitions.

When you install the Virtual I/O Server on a managed system and there is no HMC attached to the managed system when you install the Virtual I/O Server, then the Virtual I/O Server logical partition becomes the management partition. The management partition provides the Integrated Virtualization Manager Web-based system management interface and a command-line interface that you can use to manage the system. For information about using the Integrated Virtualization Manager, see Chapter 4, "Integrated Virtualization Manager," on page 163.

## Virtual I/O Server overview

Learn the concepts of the Virtual I/O Server and its primary components.

The Virtual I/O Server is software that is located in a logical partition. This software facilitates the sharing of physical I/O resources between client logical partitions within the server. The Virtual I/O Server provides virtual SCSI target and Shared Ethernet Adapter capability to client logical partitions within the system, allowing the client logical partitions to share SCSI devices and Ethernet adapters. The Virtual I/O Server software requires that the logical partition be dedicated solely for its use.

The Virtual I/O Server is available as part of the PowerVM Editions (or Advanced POWER Virtualization) hardware feature.

Using the Virtual I/O Server facilitates the following functions:

- Sharing of physical resources between logical partitions on the system
- Creating logical partitions without requiring additional physical I/O resources
- Creating more logical partitions than there are I/O slots or physical devices available with the ability for logical partitions to have dedicated I/O, virtual I/O, or both
- Maximizing use of physical resources on the system
- Helping to reduce the Storage Area Network (SAN) infrastructure

The Virtual I/O Server supports client logical partitions running the following operating systems on the following POWER6 processor-based servers.

*Table 3. Operating system support for Virtual I/O Server client logical partitions*

| **Operating system** | **POWER6 processor-based servers** |
|---|---|
| AIX 5.3 or later | All POWER6 processor-based servers |
| SUSE Linux Enterprise Server 10 Service Pack 2 or later | • 19F/HA<br>• 25F/2A |

*Table 3. Operating system support for Virtual I/O Server client logical partitions (continued)*

| Operating system | POWER6 processor-based servers |
|---|---|
| SUSE Linux Enterprise Server 10 Service Pack 1 | • 03E/4A<br>• 04E/8A<br>• 17M/MA |
| Red Hat® Enterprise Linux version 5.2 | • 19F/HA<br>• 25F/2A |
| Red Hat Enterprise Linux version 5.1 | • 03E/4A<br>• 04E/8A<br>• 17M/MA |
| Red Hat Enterprise Linux version 4.7 | 19F/HA |
| Red Hat Enterprise Linux version 4.6 | 25F/2A |
| Red Hat Enterprise Linux version 4.5 | • 03E/4A<br>• 04E/8A<br>• 17M/MA |

The Virtual I/O Server supports client logical partitions that run the following operating systems on POWER5 processor-based servers:

- AIX 5.3 (or later)
- SUSE Linux Enterprise Server 9 (or later)
- SUSE Linux Enterprise Server 10 (or later)
- Red Hat Enterprise Linux version 4 (or later)
- Red Hat Enterprise Linux version 5 (or later)

The Virtual I/O Server comprises the following primary components:

- Virtual SCSI
- Virtual networking
- Integrated Virtualization Manager

The following sections provide a brief overview of each of these components.

## Virtual SCSI

Physical adapters with attached disks or optical devices on the Virtual I/O Server logical partition can be shared by one or more client logical partitions. The Virtual I/O Server offers a local storage subsystem that provides standard SCSI-compliant logical unit numbers (LUNs). The Virtual I/O Server can export a pool of heterogeneous physical storage as a homogeneous pool of block storage in the form of SCSI disks.

Unlike typical storage subsystems that are physically located in the SAN, the SCSI devices exported by the Virtual I/O Server are limited to the domain within the server. Although the SCSI LUNs are SCSI compliant, they might not meet the needs of all applications, particularly those that exist in a distributed environment.

The following SCSI peripheral-device types are supported:

- Disk backed by a logical volume
- Disk backed by a physical volume
- Disk backed by a file

- Optical devices (DVD-RAM and DVD-ROM)
- Optical devices backed by files

## Virtual networking

Virtual I/O Server provides the following virtual networking technologies.

*Table 4. Virtual networking technologies on the Virtual I/O Server*

| Virtual networking technology | Description |
|---|---|
| Shared Ethernet Adapter | A Shared Ethernet Adapter is a layer-2 Ethernet bridge that connects physical and virtual networks together. It allows logical partitions on the virtual local area network (VLAN) to share access to a physical Ethernet adapter and to communicate with systems outside the server. Using a Shared Ethernet Adapter, logical partitions on the internal VLAN can share the VLAN with stand-alone servers.<br>On POWER6 processor-based systems, you can assign a Logical Host Ethernet port, of a Logical Host Ethernet Adapter, which is sometimes referred to as Integrated Virtual Ethernet, as the real adapter of a Shared Ethernet Adapter. A Host Ethernet Adapter is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. Host Ethernet Adapters offer high throughput, low latency, and virtualization support for Ethernet connections.<br>The Shared Ethernet Adapter on the Virtual I/O Server supports IPv6. IPv6 is the next generation of Internet protocol and is gradually replacing the current Internet standard, Internet Protocol version 4 (IPv4). The key IPv6 enhancement is the expansion of the IP address space from 32 bits to 128 bits, providing virtually unlimited, unique IP addresses. |
| Shared Ethernet Adapter failover | Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server logical partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption. |
| Link Aggregation (or EtherChannel) | A Link Aggregation (or EtherChannel) device is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters can then act as a single Ethernet device. Link Aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter. |
| Virtual local area networks (VLAN) | VLAN allows the physical network to be logically segmented. |

## Integrated Virtualization Manager

The Integrated Virtualization Manager provides a browser-based interface and a command-line interface that you can use to manage some servers that use the Virtual I/O Server. On the managed system, you can create logical partitions, manage the virtual storage and virtual Ethernet, and view service information related to the server. The Integrated Virtualization Manager is packaged with the Virtual I/O

Server, but it is activated and usable only on certain platforms and where no Hardware Management Console (HMC) is present.

## Virtual SCSI

Virtual SCSI allows client logical partitions to share disk storage and optical devices that are assigned to the Virtual I/O Server logical partition.

Disks and optical devices attached to physical adapters in the Virtual I/O Server logical partition can be shared by one or more client logical partitions. The Virtual I/O Server is a standard storage subsystem that provides standard SCSI-compliant LUNs. The Virtual I/O Server is capable of exporting a pool of heterogeneous physical storage as a homogeneous pool of block storage in the form of SCSI disks. The Virtual I/O Server is a localized storage subsystem. Unlike typical storage subsystems that are physically located in the SAN, the SCSI devices exported by the Virtual I/O Server are limited to the domain within the server. Therefore, although the SCSI LUNs are SCSI compliant, they might not meet the needs of all applications, particularly those that exist in a distributed environment.

The following SCSI peripheral device types are supported:

- Disk backed by logical volume
- Disk backed by physical volume
- Disk backed by file
- Optical CD-ROM, DVD-RAM, and DVD-ROM
- Optical DVD-RAM backed by file

Virtual SCSI is based on a client-server relationship. The Virtual I/O Server owns the physical resources as well as the *virtual SCSI server adapter*, and acts as a server, or SCSI target device. The client logical partitions have a SCSI initiator referred to as the *virtual SCSI client adapter*, and access the virtual SCSI targets as standard SCSI LUNs. You configure the virtual adapters by using the HMC or Integrated Virtualization Manager. The configuration and provisioning of virtual disk resources is performed by using the Virtual I/O Server. Physical disks owned by the Virtual I/O Server can be either exported and assigned to a client logical partition as a whole or can be partitioned into parts, such as logical volumes or files. The logical volumes and files can then be assigned to different logical partitions. Therefore, using virtual SCSI, you can share adapters as well as disk devices. To make a physical volume, logical volume, or file available to a client logical partition requires that it be assigned to a virtual SCSI server adapter on the Virtual I/O Server. The client logical partition accesses its assigned disks through a virtual-SCSI client adapter. The virtual-SCSI client adapter recognizes standard SCSI devices and LUNs through this virtual adapter.

The following figure shows a standard virtual SCSI configuration.

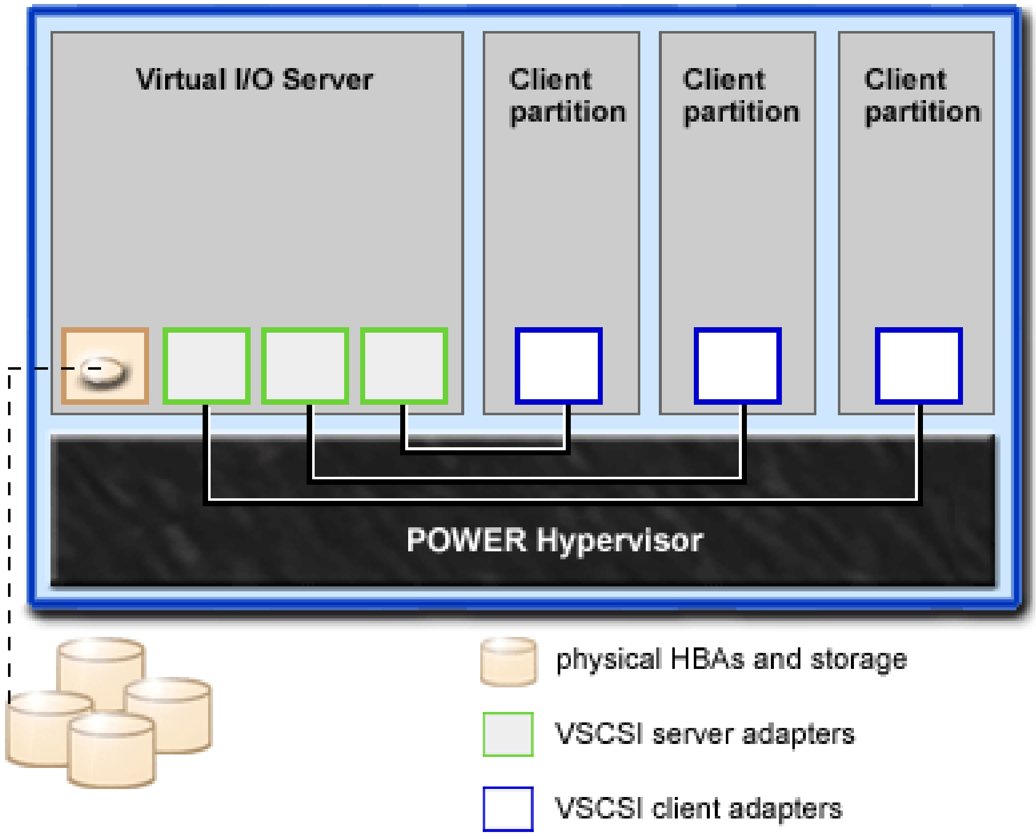

**Note:** In order for client logical partitions to be able to access virtual devices, the Virtual I/O Server must be fully operational.

## Virtual I/O Server storage subsystem overview

Learn about the Virtual I/O Server storage subsystem.

The Virtual I/O Server storage subsystem is a standard storage subsystem that provides standard SCSI-compliant LUNs. The Virtual I/O Server is a localized storage subsystem. Unlike typical storage subsystems that are physically located in the SAN, the SCSI devices exported by the Virtual I/O Server are limited to the domain within the server.

Like typical disk storage subsystems, the Virtual I/O Server has a distinct front end and back end. The front end is the interface to which client logical partitions attach to view standard SCSI-compliant LUNs. Devices on the front end are called *virtual SCSI devices*. The back end is made up of physical storage resources. These physical resources include physical disk storage, both SAN devices and internal storage devices, optical devices, logical volumes, and files.

To create a virtual device, some physical storage must be allocated and assigned to a virtual SCSI server adapter. This process creates a virtual device instance (vtscsi*X* or vtopt*X*). The device instance can be considered a mapping device. It is not a real device, but rather a mechanism for managing the mapping of the portion of physical back-end storage to the front-end virtual SCSI device. This mapping device is instrumental in re-creating the physical-to-virtual allocations in a persistent manner when the Virtual I/O Server is restarted.

## Physical storage

Learn more about physical storage, logical volumes, and the devices and configurations that are supported by the Virtual I/O Server.

**Physical volumes:**

Physical volumes can be exported to client partitions as virtual SCSI disks. The Virtual I/O Server is capable of taking a pool of heterogeneous physical disk storage attached to its back end and exporting this as homogeneous storage in the form of SCSI disk LUNs.

The Virtual I/O Server must be able to accurately identify a physical volume each time it boots, even if an event such as a storage area network (SAN) reconfiguration or adapter change has taken place. Physical volume attributes, such as the name, address, and location, might change after the system reboots due to SAN reconfiguration. However, the Virtual I/O Server must be able to recognize that this is the same device and update the virtual device mappings. For this reason, in order to export a physical volume as a virtual device, the physical volume must have either a unique identifier (UDID), a physical identifier (PVID), or an IEEE volume attribute.

For instructions on how to determine whether your disks have one of these identifiers, see "Identifying exportable disks" on page 83.

The following commands are used to manage physical volumes.

*Table 5. Physical volume commands and their descriptions*

| Physical volume command | Description |
|---|---|
| lspv | Displays information about a physical volume within a volume group. |
| migratepv | Moves allocated physical partitions from one physical volume to one or more other physical volumes. |

**Logical volumes:**

Understand how logical volumes can be exported to client partitions as virtual SCSI disks. A logical volume is a portion of a physical volume.

A hierarchy of structures is used to manage disk storage. Each individual disk drive or LUN, called a *physical volume*, has a name, such as **/dev/hdisk0**. Every physical volume in use either belongs to a volume group or is used directly for virtual storage. All of the physical volumes in a volume group are divided into physical partitions of the same size. The number of physical partitions in each region varies, depending on the total capacity of the disk drive.

Within each volume group, one or more logical volumes are defined. Logical volumes are groups of information located on physical volumes. Data on logical volumes appears to the user to be contiguous but can be discontiguous on the physical volume. This allows logical volumes to be resized or relocated and to have their contents replicated.

Each logical volume consists of one or more logical partitions. Each logical partition corresponds to at least one physical partition. Although the logical partitions are numbered consecutively, the underlying physical partitions are not necessarily consecutive or contiguous.

After installation, the system has one volume group (the rootvg volume group) consisting of a base set of logical volumes required to start the system.

You can use the commands described in the following table to manage logical volumes.

*Table 6. Logical volume commands and their descriptions*

| Logical volume command | Description |
|---|---|
| chlv | Changes the characteristics of a logical volume. |
| cplv | Copies the contents of a logical volume to a new logical volume. |
| extendlv | Increases the size of a logical volume. |
| lslv | Displays information about the logical volume. |
| mklv | Creates a logical volume. |
| mklvcopy | Creates a copy of a logical volume. |
| rmlv | Removes logical volumes from a volume group. |
| rmlvcopy | Removes a copy of a logical volume. |

Creating one or more distinct volume groups rather than using logical volumes that are created in the rootvg volume group allows you to install any newer versions of the Virtual I/O Server while maintaining client data by exporting and importing the volume groups created for virtual I/O.

**Notes:**

- Logical volumes used as virtual disks must be less than 1 TB (where TB equals 1 099 511 627 776 bytes) in size.
- For best performance, avoid using logical volumes (on the Virtual I/O Server) as virtual disks that are mirrored or striped across multiple physical volumes.

*Volume groups:*

Find information about volume groups.

A volume group is a type of storage pool that contains one or more physical volumes of varying sizes and types. A physical volume can belong to only one volume group per system. There can be up to 4096 active volume groups on the Virtual I/O Server.

When a physical volume is assigned to a volume group, the physical blocks of storage media on it are organized into physical partitions of a size determined by the system when you create the volume group. For more information, see “Physical partitions” on page 20.

When you install the Virtual I/O Server, the root volume group called rootvg is automatically created that contains the base set of logical volumes required to start the system logical partition. The rootvg includes paging space, the journal log, boot data, and dump storage, each in its own separate logical volume. The rootvg has attributes that differ from user-defined volume groups. For example, the rootvg cannot be imported or exported. When using a command or procedure on the rootvg, you must be familiar with its unique characteristics.

*Table 7. Frequently used volume group commands and their descriptions*

| Command | Description |
|---|---|
| activatevg | Activates a volume group |
| chvg | Changes the attributes of a volume group |
| deactivatevg | Deactivates a volume group |
| exportvg | Exports the definition of a volume group |
| extendvg | Adds a physical volume to a volume group |

*Table 7. Frequently used volume group commands and their descriptions (continued)*

| Command | Description |
|---|---|
| importvg | Imports a new volume group definition |
| lsvg | Displays information about a volume group |
| mkvg | Creates a volume group |
| reducevg | Removes a physical volume from a volume group |
| syncvg | Synchronizes logical volume copies that are not current |

Small systems might require only one volume group to contain all of the physical volumes (beyond the rootvg volume group). You can create separate volume groups to make maintenance easier because groups other than the one being serviced can remain active. Because the rootvg must always be online, it contains only the minimum number of physical volumes necessary for system operation. It is recommended that the rootvg not be used for client data.

You can move data from one physical volume to other physical volumes in the same volume group by using the migratepv command. This command allows you to free a physical volume so it can be removed from the volume group. For example, you could move data from a physical volume that is to be replaced.

*Physical partitions:*

This topic contains information about physical partitions.

When you add a physical volume to a volume group, the physical volume is partitioned into contiguous, equal-sized units of space called *physical partitions*. A physical partition is the smallest unit of storage space allocation and is a contiguous space on a physical volume.

Physical volumes inherit the volume group's physical partition size.

*Logical partitions:*

This topic contains information logical storage partitions.

When you create a logical volume, you specify its size in megabytes or gigabytes. The system allocates the number of logical partitions that are required to create a logical volume of at least the specified size. A logical partition is one or two physical partitions, depending on whether the logical volume is defined with mirroring enabled. If mirroring is disabled, there is only one copy of the logical volume (the default). In this case, there is a direct mapping of one logical partition to one physical partition. Each instance, including the first, is called a copy.

*Quorums:*

Find information about quorums.

A quorum exists when a majority of Volume Group Descriptor Areas and Volume Group Status Areas (VGDA/VGSA) and their disks are active. A quorum ensures data integrity of the VGDA/VGSA in the event of a disk failure. Each physical disk in a volume group has at least one VGDA/VGSA. When a volume group is created onto a single disk, the volume group initially has two VGDA/VGSA on the disk. If a volume group consists of two disks, one disk still has two VGDA/VGSA, but the other disk has one VGDA/VGSA. When the volume group is made up of three or more disks, each disk is allocated just one VGDA/VGSA.

A quorum is lost when enough disks and their VGDA/VGSA are unreachable so that a 51% majority of VGDA/VGSA no longer exists.

When a quorum is lost, the volume group deactivates itself so that the disks are no longer accessible by the logical volume manager. This prevents further disk I/O to that volume group so that data is not lost or assumed to be written when physical problems occur. As a result of the deactivation, the user is notified in the error log that a hardware error has occurred and service must be performed.

A volume group that has been deactivated because its quorum has been lost can be reactivated by using the activatevg -f command.

**Virtual media repository:**

The virtual media repository provides a single container to store and manage file-backed virtual optical media files. Media stored in the repository can be loaded into file-backed virtual optical devices for exporting to client partitions.

Only one repository can be created within a Virtual I/O Server.

The virtual media repository is available with Virtual I/O Server version 1.5 or later.

The virtual media repository is created and managed using the following commands.

*Table 8. Virtual media repository commands and their descriptions*

| Command | Description |
|---|---|
| chrep | Changes the characteristics of the virtual media repository |
| chvopt | Changes the characteristics of a virtual optical media |
| loadopt | Loads file-backed virtual optical media into a file-backed virtual optical device |
| lsrep | Displays information about the virtual media repository |
| lsvopt | Displays information about file-backed virtual optical devices |
| mkrep | Creates the virtual media repository |
| mkvdev | Creates file-backed virtual optical devices |
| mkvopt | Creates file-backed virtual optical media |
| rmrep | Removes the virtual media repository |
| rmvopt | Removes file-backed virtual optical media |
| unloadopt | Unloads file-backed virtual optical media from a file-backed virtual optical device |

**Storage pools:**

Learn about logical volume storage pools and file storage pools.

In Virtual I/O Server version 1.5 and later, you can create the following types of storage pools:
- Logical volume storage pools (LVPOOL)
- File storage pools (FBPOOL)

Like volume groups, logical volume storage pools are collections of one or more physical volumes. The physical volumes that comprise a logical volume storage pool can be of varying sizes and types. File storage pools are created within a parent logical volume storage pool and contain a logical volume containing a file system with files.

Logical volume storage pools store logical volume backing devices, file-backed storage pools, and the virtual media repository. File storage pools store file-backing devices.

Using storage pools, you are not required to have extensive knowledge of how to manage volume groups and logical volumes to create and assign logical storage to a client logical partition. Devices created using a storage pool are not limited to the size of the individual physical volumes.

Storage pools are created and managed using the following commands.

*Table 9. Storage pool commands and their descriptions*

| Command | Description |
|---|---|
| chsp | Changes the characteristics of a storage pool |
| chbdsp | Changes the characteristics of a backing device within a storage pool |
| lssp | Displays information about a storage pool |
| mkbdsp | Assigns storage from a storage pool to be a backing device for a virtual SCSI adapter |
| mksp | Creates a storage pool |
| rmdbsp | Disassociates a backing device from its virtual SCSI adapter and removes it from the system |
| rmsp | Removes a file storage pool |

Each Virtual I/O Server logical partition has a single default storage pool that can be modified only by the prime administrator. If the default storage pool is not modified by the prime administrator, rootvg, which is a logical volume pool, is used as the default storage pool.

Do not create client storage in rootvg. Creating one or more distinct logical volume storage pools rather than using the rootvg volume group allows you to install any newer versions of the Virtual I/O Server while maintaining client data by exporting and importing the volume groups created for virtual I/O.

Unless explicitly specified otherwise, the storage pool commands will operate on the default storage pool. This situation can be useful on systems that contain most or all of its backing devices in a single storage pool.

**Note:** Storage pools cannot be used when assigning whole physical volumes as backing devices.

**Optical devices:**

Optical devices can be exported by the Virtual I/O Server. This topic gives information about what types of optical devices are supported.

The Virtual I/O Server supports exporting optical SCSI devices. These are referred to as a *virtual SCSI optical devices*. Virtual optical devices can be backed by DVD drives or files. Depending on the backing device, the Virtual I/O Server will export a virtual optical device with one of following profiles:

- DVD-ROM
- DVD-RAM

Virtual optical devices that are backed by physical optical devices can be assigned to only one client logical partition at a time. In order to use the device on a different client logical partition, it must first be removed from its current logical partition and reassigned to the logical partition that will use the device.

## Virtual storage

Disks and optical devices are supported as virtual SCSI devices. This topic describes how those devices function in a virtualized environment and provides information on what devices are supported.

**Disk:**

Disk devices can be exported by the Virtual I/O Server. This topic gives information about what types of disks and configurations are supported.

The Virtual I/O Server supports exporting disk SCSI devices. These are referred to as *virtual SCSI disks*. All virtual SCSI disks must be backed by physical storage. The following types of physical storage can be used to back virtual disks:

- Virtual SCSI disk backed by a physical disk
- Virtual SCSI disk backed by a logical volume
- Virtual SCSI disk backed by a file

Regardless of whether the virtual SCSI disk is backed by a physical disk, logical volume, or a file, all standard SCSI rules apply to the device. The virtual SCSI device will behave as a standard SCSI-compliant disk device, and it can serve as a boot device or a Network Installation Management (NIM) target, for example.

**Virtual SCSI Client Adapter Path Timeout**

The virtual SCSI (VSCSI) Client Adapter Path Timeout feature allows the client adapter to detect whether a Virtual I/O Server is not responding to I/O requests. Use this feature only in configurations in which devices are available to a client logical partition from multiple Virtual I/O Servers. These configurations could be either configurations where Multipath I/O (MPIO) is being used or where a volume group is being mirrored by devices on multiple Virtual I/O Servers.

If no I/O requests issued to the VSCSI server adapter have been serviced within the number of seconds specified by the VSCSI path timeout value, one more attempt is made to contact the VSCSI server adapter, waiting up to 60 seconds for a response.

If, after 60 seconds, there is still no response from the server adapter, all outstanding I/O requests to that adapter are failed and an error is written to the client logical partition error log. If MPIO is being used, the MPIO Path Control Module will retry the I/O requests down another path. Otherwise, the failed requests will be returned to the applications. If the devices on this adapter are part of a mirrored volume group, those devices will be marked as *missing* and the Logical Volume Manager logs errors in the client logical partition error log. If one of the failed devices is the root volume group (rootvg) for the logical partition, and the rootvg is not available via another path or is not being mirrored on another Virtual I/O Server, the client logical partition is likely to shut down. The VSCSI client adapter attempts to reestablish communication with the Virtual I/O Server and logs a message in the system error log when it is able to do so. Mirrored volume groups must be manually resynchronized by running the varyonvg command when the missing devices are once again available.

A configurable VSCSI client adapter ODM attribute, **vscsi_path_to**, is provided. This attribute is used to both indicate if the feature is enabled and to store the value of the path timeout if the feature is enabled.

The system administrator sets the ODM attribute to 0 to disable the feature, or to the time, in seconds, to wait before checking if the path to the server adapter has failed. If the feature is enabled, a minimum setting of 30 seconds is required. If a setting between 0 and 30 seconds is entered, the value will be changed to 30 seconds upon the next adapter reconfiguration or reboot.

This feature is disabled by default, thus the default value of **vscsi_path_to** is 0. Exercise careful consideration when setting this value, keeping in mind that when the VSCSI server adapter is servicing the I/O request, the storage device the request is being sent to may be either local to the VIO Server or on a SAN.

The **vscsi_path_to** client adapter attribute can be set by using the SMIT utility or by using the **chdev -P** command. The attribute setting can also be viewed by using SMIT or the **lsattr** command. The setting will not take affect until the adapter is reconfigured or the machine is rebooted.

**Optical:**

Optical devices can be exported by the Virtual I/O Server. This topic gives information about what types of optical devices are supported.

The Virtual I/O Server supports exporting physical optical devices to client logical partitions. These are referred to as *virtual SCSI optical devices*. Virtual SCSI optical devices can be backed by DVD drives or files. Depending on the backing device, the Virtual I/O Server will export a virtual optical device with one of following profiles:

- DVD-ROM
- DVD-RAM

For example, file-backed virtual SCSI optical devices are exported as DVD-RAM devices. File-backed virtual SCSI optical devices can be backed by read-write or read-only files. Depending on the file permissions, the device can appear to contain a DVD-ROM or DVD-RAM disk. Read-write media files (DVD-RAM) cannot be loaded into more than one file-backed virtual SCSI optical device simultaneously. Read-only media files (DVD-ROM) can be loaded into multiple file-backed virtual SCSI optical devices simultaneously.

Virtual SCSI optical devices that are backed by physical optical devices can be assigned to only one client logical partition at any given time. To use the device on a different client logical partition, it must first be removed from its current logical partition and reassigned to the logical partition that will use the device.

Virtual SCSI optical devices will always appear as SCSI devices on the client logical partitions regardless of whether the device type exported from the Virtual I/O Server is a SCSI, IDE, USB device, or a file.

### Mapping devices

Mapping devices are used to facilitate the mapping of physical resources to a virtual device.

## Virtual networking

Learn about virtual Ethernet, Host Ethernet Adapter (or Integrated Virtual Ethernet), Internet Protocol version 6 (IPv6), Link Aggregation (or EtherChannel), Shared Ethernet Adapter, Shared Ethernet Adapter failover, and VLAN.

Virtual Ethernet technology facilitates IP-based communication between logical partitions on the same system using virtual local area network (VLAN)-capable software switch systems. Using Shared Ethernet Adapter technology, logical partitions can communicate with other systems outside the hardware unit without assigning physical Ethernet slots to the logical partitions.

### Host Ethernet Adapter

A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, multiple logical partitions can connect directly to the HEA and use the HEA resources. This allows these logical partitions to access external networks through the HEA without having to go through an Ethernet bridge on another logical partition.

To connect a logical partition to an HEA, you must create a Logical Host Ethernet Adapter (LHEA) for the logical partition. A *Logical Host Ethernet Adapter (LHEA)* is a representation of a physical HEA on a logical partition. An LHEA appears to the operating system as if it were a physical Ethernet adapter, just as a virtual Ethernet adapter appears as if it were a physical Ethernet adapter. When you create an LHEA for a logical partition, you specify the resources that the logical partition can use on the actual physical

HEA. Each logical partition can have one LHEA for each physical HEA on the managed system. Each LHEA can have one or more logical ports, and each logical port can connect to a physical port on the HEA.

You can create an LHEA for a logical partition using either of the following methods:

- You can add the LHEA to a partition profile, shut down the logical partition, and reactivate the logical partition using the partition profile with the LHEA.
- You can add the LHEA to a running logical partition using dynamic logical partitioning. (This method can be used for Linux logical partitions only if you install Red Hat Enterprise Linux version 5.1, Red Hat Enterprise Linux version 4.6, or a later version of Red Hat Enterprise Linux on the logical partition.)

When you activate a logical partition, the LHEAs in the partition profile are considered to be required resources. If the physical HEA resources required by the LHEAs are not available, then the logical partition cannot be activated. However, when the logical partition is active, you can remove any LHEAs you want from the logical partition.

After you create an LHEA for a logical partition, a network device is created in the logical partition. This network device is named `entX` on AIX logical partitions and `ethX` on Linux logical partitions, where `X` represents sequentially assigned numbers. The user can then set up TCP/IP configuration similar to a physical Ethernet device to communicate with other logical partitions.

A logical port can communicate with all other logical ports that are connected to the same physical port on the HEA. The physical port and its associated logical ports form a logical Ethernet network. Broadcast and multicast packets are distributed on this logical network as though it was a physical Ethernet network. You can connect up to 16 logical ports to a physical port using this logical network. By extension, you can connect up to 16 logical partitions to each other and to an external network through this logical network. The actual number of logical ports that you can connect to a physical port depends upon the Multi-Core Scaling value of the physical port group and the number of logical ports that have been created for other physical ports within the physical port group. By default, the Multi-Core Scaling value of each physical port group is set to 4, which allows 4 logical ports to be connected to the physical ports in the physical port group. To allow up to 16 logical ports to be connected to the physical ports in the physical port group, you must change the Multi-Core Scaling value of the physical port group to 1 and restart the managed system.

If you want to connect more than 16 logical partitions to each other and to an external network through a physical port on an HEA, you can create a logical port on a Virtual I/O Server logical partition and configure an Ethernet bridge between the logical port and a virtual Ethernet adapter on a virtual LAN. This allows all logical partitions with virtual Ethernet adapters on the virtual LAN to communicate with the physical port through the Ethernet bridge. If you configure an Ethernet bridge between a logical port and a virtual Ethernet adapter, the physical port that is connected to the logical port must have the following properties:

- The physical port must be configured so that the Virtual I/O Server logical partition is the promiscuous mode partition for the physical port. For more information on how to configure a physical port, see Configuring physical ports on a Host Ethernet Adapter using the HMC.
- The physical port can have only one logical port.

You can set each logical port to restrict or allow packets that are tagged for specific VLANs. You can set a logical port to accept packets with any VLAN ID, or you can set a logical port to accept only the VLAN IDs that you specify. You can specify up to 20 individual VLAN IDs for each logical port.

The physical ports on an HEA are always configured on the managed system level. If you use an HMC to manage a system, you must use the HMC to configure the physical ports on any HEAs belonging to the managed system. Also, the physical port configuration applies to all logical partitions that use the physical port. (Some properties might require setup in the operating system as well. For example, the maximum packet size for a physical port on the HEA must be set on the managed system level using the HMC. However, you must also set the maximum packet size for each logical port within the operating system.) By contrast, if a system is unpartitioned and is not managed by an HMC, you can configure the physical ports on an HEA within the operating system just as if the physical ports were ports on a regular physical Ethernet adapter.

HEA hardware does not support Half Duplex mode.

You can change the properties of a logical port on an LHEA by using dynamic logical partitioning to remove the logical port from the logical partition and add the logical port back to the logical partition using the changed properties. If the operating system of the logical partition does not support dynamic logical partitioning for LHEAs, and you want to change any logical port property other than the VLANs on which the logical port participates, you must set a partition profile for the logical partition so that the partition profile contains the desired logical port properties, shut down the logical partition, and activate the logical partition using the new or changed partition profile. If the operating system of the logical partition does not support dynamic logical partitioning for LHEAs, and you want to change the VLANs on which the logical port participates, you must remove the logical port from a partition profile belonging to the logical partition, shut down and activate the logical partition using the changed partition profile, add the logical port back to the partition profile using the changed VLAN configuration, and shut down and activate the logical partition again using the changed partition profile.

### Internet Protocol version 6

Internet Protocol version 6 (IPv6) is the next generation of Internet Protocol and is gradually replacing the current Internet standard, Internet Protocol version 4 (IPv4). The key IPv6 enhancement is the expansion of the IP address space from 32 bits to 128 bits, providing virtually unlimited, unique IP addresses.

IPv6 provides several advantages over IPv4, including expanded routing and addressing, routing simplification, header format simplification, improved traffic control, autoconfiguration, and security.

For more information about IPv6, see Internet Protocol (IP) Version 6.

### Link Aggregation or EtherChannel devices

A Link Aggregation, or EtherChannel, device is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters can then act as a single Ethernet device. Link Aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

For example, `ent0` and `ent1` can be aggregated to `ent3`. The system considers these aggregated adapters as one adapter, and all adapters in the Link Aggregation device are given the same hardware address, so they are treated by remote systems as if they are one adapter.

Link Aggregation can help provide more redundancy because individual links might fail, and the Link Aggregation device will fail over to another adapter in the device to maintain connectivity. For example, in the previous example, if `ent0` fails, the packets are automatically sent on the next available adapter, `ent1`, without disruption to existing user connections. `ent0` automatically returns to service on the Link Aggregation device when it recovers.

You can configure a Shared Ethernet Adapter to use a Link Aggregation, or EtherChannel, device as the physical adapter.

### Virtual Ethernet adapters

Virtual Ethernet adapters allow client logical partitions to send and receive network traffic without having a physical Ethernet adapter.

Virtual Ethernet adapters allow logical partitions within the same system to communicate without having to use physical Ethernet adapters. Within the system, virtual Ethernet adapters are connected to an IEEE

802.1q virtual Ethernet switch. Using this switch function, logical partitions can communicate with each other by using virtual Ethernet adapters and assigning VIDs. With VIDs, virtual Ethernet adapters can share a common logical network. The system transmits packets by copying the packet directly from the memory of the sender logical partition to the receive buffers of the receiver logical partition without any intermediate buffering of the packet.

Virtual Ethernet adapters can be used without using the Virtual I/O Server, but the logical partitions will not be able to communicate with external systems. However, in this situation, you can use another device, called a Host Ethernet Adapter (or Integrated Virtual Ethernet), to facilitate communication between logical partitions on the system and external networks.

You can create virtual Ethernet adapters using the Hardware Management Console (HMC) and configure them using the Virtual I/O Server command-line interface. You can also use the Integrated Virtualization Manager to create and manage virtual Ethernet adapters.

Consider using virtual Ethernet on the Virtual I/O Server in the following situations:

- When the capacity or the bandwidth requirement of the individual logical partition is inconsistent with, or is less than, the total bandwidth of a physical Ethernet adapter. Logical partitions that use the full bandwidth or capacity of a physical Ethernet adapter should use dedicated Ethernet adapters.
- When you need an Ethernet connection, but there is no slot available in which to install a dedicated adapter.

## Virtual local area networks

Virtual local area networks (VLAN) allows the physical network to be logically segmented.

VLAN is a method to logically segment a physical network so that layer 2 connectivity is restricted to members that belong to the same VLAN. This separation is achieved by tagging Ethernet packets with their VLAN membership information and then restricting delivery to members of that VLAN. VLAN is described by the IEEE 802.1Q standard.

The VLAN tag information is referred to as VLAN ID (VID). Ports on a switch are configured as being members of a VLAN designated by the VID for that port. The default VID for a port is referred to as the Port VID (PVID). The VID can be added to an Ethernet packet either by a VLAN-aware host, or by the switch in the case of VLAN-unaware hosts. Ports on an Ethernet switch must therefore be configured with information indicating whether the host connected is VLAN-aware.

For VLAN-unaware hosts, a port is set up as untagged and the switch will tag all packets entering through that port with the Port VLAN ID (PVID). It will also untag all packets exiting that port before delivery to the VLAN unaware host. A port used to connect VLAN-unaware hosts is called an *untagged port*, and it can be a member of only a single VLAN identified by its PVID. Hosts that are VLAN-aware can insert and remove their own tags and can be members of more than one VLAN. These hosts are typically attached to ports that do not remove the tags before delivering the packets to the host, but will insert the PVID tag when an untagged packet enters the port. A port will only allow packets that are untagged or tagged with the tag of one of the VLANs that the port belongs to. These VLAN rules are in addition to the regular media access control (MAC) address-based forwarding rules followed by a switch. Therefore, a packet with a broadcast or multicast destination MAC is also delivered to member ports that belong to the VLAN that is identified by the tags in the packet. This mechanism ensures the logical separation of the physical network based on membership in a VLAN.

## Shared Ethernet Adapters

Shared Ethernet Adapters on the Virtual I/O Server logical partition allow virtual Ethernet adapters on client logical partitions to send and receive outside network traffic.

A Shared Ethernet Adapter is a Virtual I/O Server component that bridges a physical Ethernet adapter and one or more virtual Ethernet adapters:

- The real adapter can be a physical Ethernet adapter, a Link Aggregation or EtherChannel device, or a Logical Host Ethernet Adapter. The real adapter cannot be another Shared Ethernet Adapter, a VLAN pseudo-device, or a virtual Ethernet adapter.
- The virtual Ethernet adapter must be a virtual Ethernet adapter. It cannot be any other type of device or adapter.

Using a Shared Ethernet Adapter, logical partitions on the virtual network can share access to the physical network and communicate with stand-alone servers and logical partitions on other systems. The Shared Ethernet Adapter eliminates the need for each client logical partition to own a real adapter to connect to the external network.

A Shared Ethernet Adapter provides access by connecting the internal VLANs with the VLANs on the external switches. Using this connection, logical partitions can share the IP subnet with stand-alone systems and other external logical partitions. The Shared Ethernet Adapter forwards outbound packets received from a virtual Ethernet adapter to the external network and forwards inbound packets to the appropriate client logical partition over the virtual Ethernet link to that logical partition. The Shared Ethernet Adapter processes packets at layer 2, so the original MAC address and VLAN tags of the packet are visible to other systems on the physical network.

### GARP VLAN Registration Protocol

Shared Ethernet Adapters, in Virtual I/O Server version 1.4 or later, support GARP VLAN Registration Protocol (GVRP), which is based on GARP (Generic Attribute Registration Protocol). GVRP allows for the dynamic registration of VLANs over networks, which can reduce the number of errors in the configuration of a large network. By propagating registration across the network through the transmission of Bridge Protocol Data Units (BPDUs), devices on the network have accurate knowledge of the bridged VLANs configured on the network.

When GVRP is enabled, communication travels one way: from the Shared Ethernet Adapter to the switch. The Shared Ethernet Adapter notifies the switch which VLANs can communicate with the network. The Shared Ethernet Adapter does not configure VLANs to communicate with the network based on information received from the switch. Rather, the configuration of VLANs to communicate with the network is statically determined by the virtual Ethernet adapter configuration settings.

### Host Ethernet Adapter or Integrated Virtual Ethernet

With Virtual I/O Server version 1.4, you can assign a Logical Host Ethernet port, of a Logical Host Ethernet Adapter (LHEA), which is sometimes referred to as Integrated Virtual Ethernet, as the real adapter of a Shared Ethernet Adapter. The Logical Host Ethernet port is associated with a physical port on the Host Ethernet Adapter. The Shared Ethernet Adapter uses the standard device driver interfaces provided by the Virtual I/O Server to interface with the Host Ethernet Adapter.

To use a Shared Ethernet Adapter with a Host Ethernet Adapter, the following requirements must be met:

- The Logical Host Ethernet port must be the only port assigned to the physical port on the Host Ethernet Adapter. No other ports of the LHEA can be assigned to the physical port on the Host Ethernet Adapter.
- The LHEA on the Virtual I/O Server logical partition must be set to promiscuous mode. (In an Integrated Virtualization Manager environment, the mode is set to *promiscuous* by default.) *Promiscuous* mode allows the LHEA (on the Virtual I/O Server) to receive all unicast, multicast, and broadcast network traffic from the physical network.

### Recommendations

Consider using Shared Ethernet Adapters on the Virtual I/O Server in the following situations:

- When the capacity or the bandwidth requirement of the individual logical partition is inconsistent or is less than the total bandwidth of a physical Ethernet adapter. Logical partitions that use the full bandwidth or capacity of a physical Ethernet adapter should use dedicated Ethernet adapters.
- If you plan to migrate a client logical partition from one system to another.

Consider assigning a Shared Ethernet Adapter to a Logical Host Ethernet port when the number of Ethernet adapters that you need is more than the number of ports available on the LHEA, or you anticipate that your needs will grow beyond that number. If the number of Ethernet adapters that you need is fewer than or equal to the number of ports available on the LHEA, and you do not anticipate needing more ports in the future, then you can use the ports of the LHEA for network connectivity rather than the Shared Ethernet Adapter.

## Virtual I/O Server management

Learn about management tools for the Virtual I/O Server, such as the Virtual I/O Server command-line interface, and several Tivoli® products that can manage different aspects of the Virtual I/O Server.

For systems that are not managed by a Hardware Management Console (HMC), the Virtual I/O Server becomes the management partition and provides a graphical user interface, called the Integrated Virtualization Manager, to help you manage the system. For more information, see Chapter 4, “Integrated Virtualization Manager,” on page 163.

### Virtual I/O Server command-line interface

Learn about accessing and using the Virtual I/O Server command-line interface.

The Virtual I/O Server is configured and managed through a command-line interface. In environments where no HMC is present, some Virtual I/O Server tasks can also be performed using the Integrated Virtualization Manager. All aspects of Virtual I/O Server administration can be accomplished through the command-line interface, including the following:
- Device management (physical, virtual, logical volume manager (LVM))
- Network configuration
- Software installation and update
- Security
- User management
- Maintenance tasks

In addition, in environments managed by the Integrated Virtualization Manager, you can use the Virtual I/O Server command-line interface to manage logical partitions.

The first time you log in to the Virtual I/O Server, use the **padmin** user ID, which is the prime administrator user ID. You will be prompted for a new password.

#### Restricted shell

Upon logging in, you will be placed into a restricted Korn shell. The restricted Korn shell works in the same way as a standard Korn shell, except that you cannot do the following:
- Change the current working directory
- Set the value of the **SHELL**, **ENV**, or **PATH** variables
- Specify the path name of the command that contains a forward slash (/)
- Redirect output of a command using any of the following characters: >, >|, <>, >>

As a result of these restrictions, you will not be able to execute commands that are not accessible to your **PATH** variables. In addition, these restrictions prevent you from sending command output directly to a file. Instead, command output can be piped to the tee command.

After you log in, you can type `help` to get information about the supported commands. For example, to get help on the errlog command, type `help errlog`.

## Execution Mode

The Virtual I/O Server command-line interface functions similarly to a standard command-line interface. Commands are issued with appropriate accompanying flags and parameters. For example, to list all adapters, type the following:

```
lsdev -type adapter
```

In addition, scripts can be run within the Virtual I/O Server command-line interface environment.

In addition to the Virtual I/O Server command-line interface commands, the following standard shell commands are provided.

*Table 10. Standard shell commands and their functions*

| **Command** | **Function** |
|---|---|
| awk | Matches patterns and performs actions on them. |
| cat | Concatenates or displays files. |
| chmod | Changes file modes. |
| cp | Copies files. |
| date | Displays the date and time. |
| grep | Searches a file for a pattern. |
| ls | Displays the contents of a directory |
| mkdir | Makes a directory. |
| man | Displays manual entries for the Virtual I/O Server commands. |
| more | Displays the contents of files one screen at a time. |
| rm | Removes files. |
| sed | Provides a stream editor. |
| stty | Sets, resets, and reports workstation operating parameters. |
| tee | Displays the output of a program and copies it to a file. |
| vi | Edits files with full screen display. |
| wc | Counts the number of lines, words, bytes, and characters in a file |
| who | Identifies the users currently logged in. |

As each command is executed, the user log and the global command log are updated.

The user log will contain a list of each Virtual I/O Server command, including arguments, that a user has executed. One user log for each user in the system is created. This log is located in the user's home directory and can be viewed by using either the cat or the vi commands.

The global command log is made up of all the Virtual I/O Server command-line interface commands executed by all users, including arguments, the date and time the command was executed, and from which user ID it was executed. The global command log is viewable only by the **padmin** user ID, and it can be viewed by using the lsgcl command. If the global command log exceeds 1 MB, the log will be truncated to 250 KB to prevent the file system from reaching capacity.

**Note:** Integrated Virtualization Manager commands are audited in a separate place and are viewable either in **Application Logs**, or by running the following command from the command line:

```
lssvcevents -t console --filter severities=audit
```

**Related information**

Virtual I/O Server and Integrated Virtualization Manager Command Reference

## Tivoli software and the Virtual I/O Server

Learn about integrating the Virtual I/O Server into your Tivoli environment for Tivoli Application Dependency Discovery Manager, Tivoli Monitoring, Tivoli Storage Manager, Tivoli Usage and Accounting Manager, Tivoli Identity Manager, and TotalStorage® Productivity Center.

### Tivoli Application Dependency Discovery Manager

Tivoli Application Dependency Discovery Manager (TADDM) discovers infrastructure elements found in the typical data center, including application software, hosts and operating environments (including the Virtual I/O Server), network components (such as routers, switches, load balancers, firewalls, and storage), and network services (such as LDAP, NFS, and DNS). Based on the data it collects, TADDM automatically creates and maintains application infrastructure maps that include runtime dependencies, configuration values, and change history. With this information, you can determine the interdependences between business applications, software applications, and physical components to help you ensure and improve application availability in your environment. For example, you can do the following tasks:

- You can isolate configuration-related application problems.
- You can plan for application changes to minimize or eliminate unplanned disruptions.
- You can create a shared topological definition of applications for use by other management applications.
- You can determine the effect of a single configuration change on a business application or service.
- You can see what changes take place in the application environment and where.

TADDM includes an agent-free discovery engine, which means that the Virtual I/O Server does not require that an agent or client be installed and configured in order to be discovered by TADDM. Instead, TADDM uses discovery sensors that rely on open and secure protocols and access mechanisms to discover the data center components.

For more information, see the Tivoli Application Dependency Discovery Manager Information Center.

### Tivoli Identity Manager

With Tivoli Identity Manager, you can manage identities and users across several platforms, including AIX, Windows®, Solaris, and so on. With Tivoli Identity Manager 4.7, you can also include Virtual I/O Server users. Tivoli Identity Manager provides a Virtual I/O Server adapter that acts as an interface between the Virtual I/O Server and the Tivoli Identity Manager Server. The adapter might not be located on the Virtual I/O Server and the Tivoli Identity Manager Server manages access to the Virtual I/O Server by using your security system. The adapter runs as a service, independent of whether a user is logged on to the Tivoli Identity Manager Server. The adapter acts as a trusted virtual administrator on the Virtual I/O Server, performing tasks like the following:

- Creating a user ID to authorize access to the Virtual I/O Server.
- Modifying an existing user ID to access the Virtual I/O Server.
- Removing access from a user ID. This deletes the user ID from the Virtual I/O Server.
- Suspending a user account by temporarily deactivating access to the Virtual I/O Server.
- Restoring a user account by reactivating access to the Virtual I/O Server.
- Changing a user account password on the Virtual I/O Server.
- Reconciling the user information of all current users on the Virtual I/O Server.
- Reconciling the user information of a particular user account on the Virtual I/O Server by performing a lookup.

For more information, see the Tivoli Identity Manager product manuals.

### Tivoli Monitoring

Virtual I/O Server V1.3.0.1 (fix pack 8.1), includes the Tivoli Monitoring System Edition agent. With Tivoli Monitoring System Edition , you can monitor the health and availability of multiple servers (including the Virtual I/O Server) from the Tivoli Enterprise™ Portal. Tivoli Monitoring System Edition gathers data from the Virtual I/O Server, including data about physical volumes, logical volumes, storage pools, storage mappings, network mappings, real memory, processor resources, mounted file system sizes, and so on. From the Tivoli Enterprise Portal, you can view a graphical representation of the data, use predefined thresholds to alert you on key metrics, and resolve issues based on recommendations provided by the Expert Advice feature of Tivoli Monitoring.

For more information, see the following resources:

- Tivoli Monitoring 6.1 documentation
- Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide

### Tivoli Storage Manager

Virtual I/O Server 1.4 includes the Tivoli Storage Manager client. With Tivoli Storage Manager, you can protect Virtual I/O Server data from failures and other errors by storing backup and disaster-recovery data in a hierarchy of offline storage. Tivoli Storage Manager can help protect computers running a variety of different operating environments, including the Virtual I/O Server, on a variety of different hardware. If you configure the Tivoli Storage Manager client on the Virtual I/O Server, you can include the Virtual I/O Server in your standard backup framework.

For more information, see Tivoli Storage Manager for UNIX® and Linux Backup-Archive Clients Installation and User's Guide

### Tivoli Usage and Accounting Manager

Virtual I/O Server 1.4 includes the Tivoli Usage and Accounting Manager agent on the Virtual I/O Server. Tivoli Usage and Accounting Manager helps you track, allocate, and invoice your IT costs by collecting, analyzing, and reporting on the actual resources used by entities such as cost centers, departments, and users. Tivoli Usage and Accounting Manager can gather data from multi-tiered datacenters that include Windows, AIX, Virtual I/O Server, HP/UX Sun Solaris, Linux, and VMware.

For more information, see the Tivoli Usage and Accounting Manager Information Center.

### TotalStorage Productivity Center

With Virtual I/O Server 1.5.2, you can install and configure the TotalStorage Productivity Center agents on the Virtual I/O Server. TotalStorage Productivity Center is an integrated, storage infrastructure management suite that is designed to help simplify and automate the management of storage devices, storage networks, and capacity utilization of file systems and databases. When you install and configure the TotalStorage Productivity Center agents on the Virtual I/O Server, you can use the TotalStorage Productivity Center user interface to collect and view information about the Virtual I/O Server. You can then perform the following tasks using the TotalStorage Productivity Center user interface:

1. Run a discovery job for the agents on the Virtual I/O Server.
2. Run probes, run scans, and ping jobs to collect storage information about the Virtual I/O Server.
3. Generate reports using the Fabric Manager and the Data Manager to view the storage information gathered.
4. View the storage information gathered using the topology Viewer.

For more information, see the *TotalStorage Productivity Center support for agents on a Virtual I/O Server* PDF. To view or download the PDF, go to the TotalStorage Productivity Center v3.3.1.81 Interim Fix Web site.

**Related tasks**

"Configuring the Tivoli agents and clients on the Virtual I/O Server" on page 88
You can configure and start the Tivoli Monitoring agent, Tivoli Usage and Accounting Manager, the Tivoli Storage Manager client, and the Tivoli TotalStorage Productivity Center agents.

## Configuration scenarios for the Virtual I/O Server

The following scenarios show examples of networking configurations for the Virtual I/O Server logical partition and the client logical partitions. Use the following scenarios and configuration examples to understand more about the Virtual I/O Server and its components.

## Scenario: Configuring a Virtual I/O Server without VLAN tagging

Use this scenario to help you become familiar with creating a network without VLAN tagging.

**Situation**

You are the system administrator responsible for planning and configuring the network in an environment with the Virtual I/O Server running. You want to configure a single logical subnet on the system that communicates with the switch.

**Objective**

The objective of this scenario is to configure the network where only Port Virtual LAN ID (PVID) is used, the packets are not tagged, and a single internal network is connected to a switch. There are no virtual local area networks (VLAN) tagged ports set up on the Ethernet switch, and all virtual Ethernet adapters are defined using a single default PVID and no additional VLAN IDs (VIDs).

**Prerequisites and assumptions**

- The Hardware Management Console (HMC) was set up. To view the PDF file of the *Installation and Configuration Guide for the Hardware Management Console* (SA76-0084), approximately 3 MB in size, see sa76-0084.pdf .
- You understand the partitioning concepts as described in the *Logical Partitioning Guide*. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf .
- The Virtual I/O Server logical partition has been created and the Virtual I/O Server has been installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 61.
- You have created the remaining logical partitions that you want added to the network configuration.
- You have an Ethernet switch and a router ready to add to the configuration.
- You have IP addresses for all logical partitions and systems that will be added to the configuration.

While this procedure describes configuration in an HMC environment, this configuration is also possible in an Integrated Virtualization Manager environment.

**Configuration steps**

The following figure shows the configuration that will be completed during this scenario.

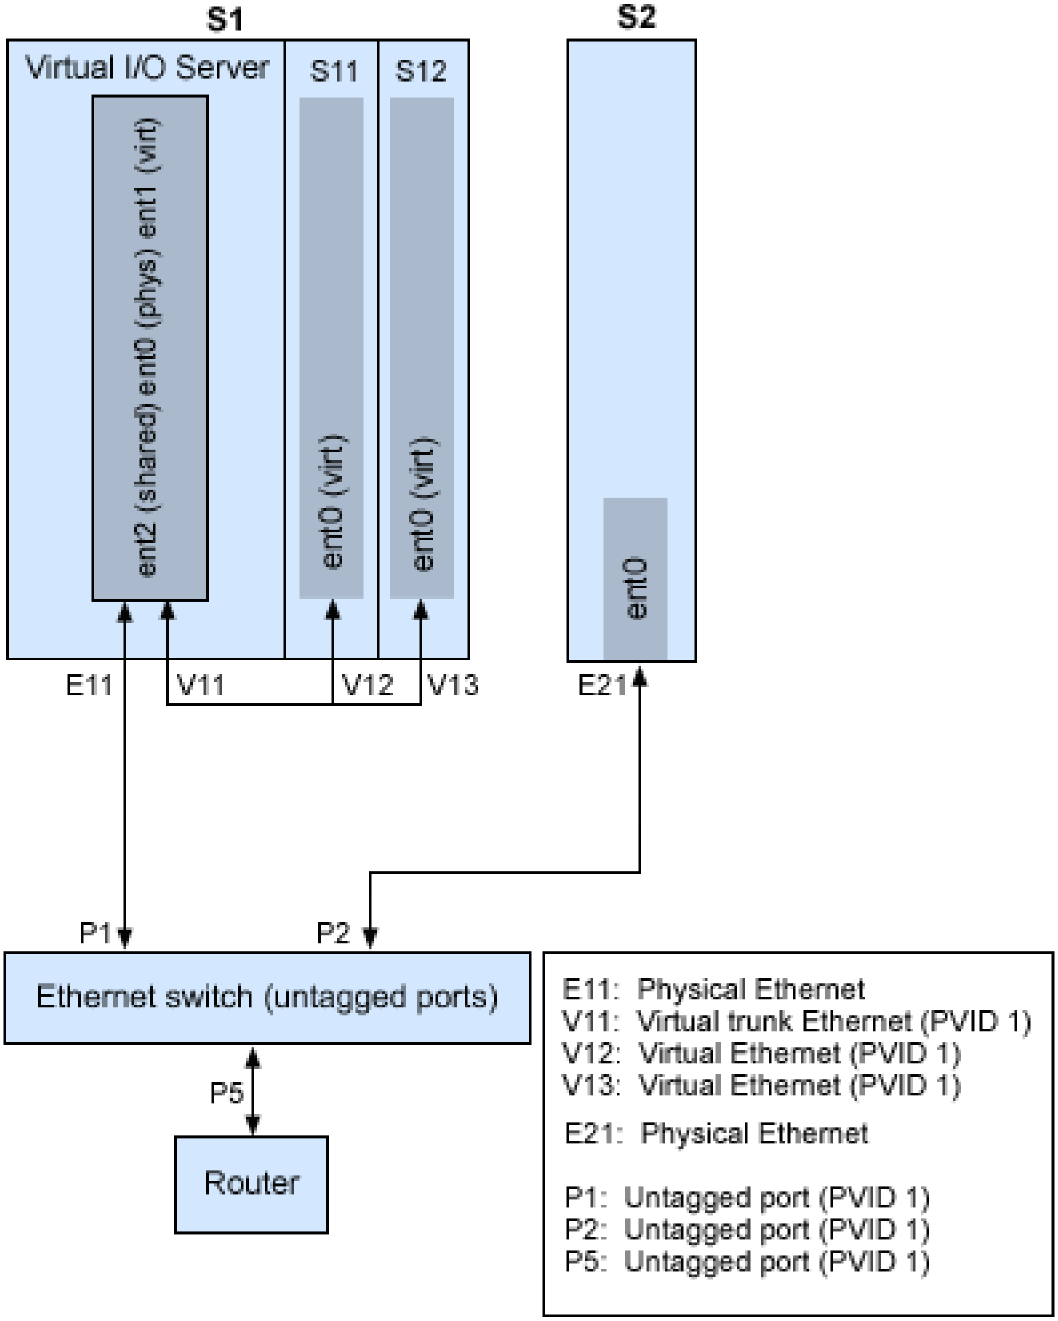
S1
S2
Virtual I/O Server
S11
S12
ent2 (shared) ent0 (phys) ent1 (virt)
ent0 (virt)
ent0 (virt)
ent0
E11
V11
V12
V13
E21
P1
P2
Ethernet switch (untagged ports)
P5
Router
E11: Physical Ethernet
V11: Virtual trunk Ethernet (PVID 1)
V12: Virtual Ethernet (PVID 1)
V13: Virtual Ethernet (PVID 1)
E21: Physical Ethernet
P1: Untagged port (PVID 1)
P2: Untagged port (PVID 1)
P5: Untagged port (PVID 1)

Using the preceding figure as a guide, follow these steps:

1. Set up an Ethernet switch with untagged ports. Alternatively, you can use an Ethernet switch that does not use VLAN.
2. For system S1, use the HMC to create a virtual Ethernet adapter (V11) for the Virtual I/O Server with the trunk setting, PVID set to 1, and no additional VIDs.
3. For system S1, use the HMC to create virtual Ethernet adapters V12 and V13 for logical partitions S11 and S12, respectively, with PVID set to 1 and no additional VIDs.
4. For system S1, use the HMC to assign physical Ethernet adapter E11 to the Virtual I/O Server and connect the adapter to the Ethernet switch port P1.

5. On the Virtual I/O Server, set up Shared Ethernet Adapter ent2 with the physical adapter ent0 and virtual adapter ent1.
6. Start the logical partitions. The process recognizes the virtual devices that were created in Step 1.
7. Configure IP addresses for S11 (en0), S12 (en0), and S2 (en0), so that they all belong to the same subnet with the router connected to Ethernet switch port P5.

The Shared Ethernet Adapter on the Virtual I/O Server logical partition can also be configured with IP addresses on the same subnet. This is required only for network connectivity to the Virtual I/O Server logical partition.

## Scenario: Configuring a Virtual I/O Server using VLAN tagging

Use this scenario to help you become familiar with creating a network using VLAN tagging.

### Situation

You are the system administrator responsible for planning and configuring the network in an environment with the Virtual I/O Server running. You would like to configure the network so that two logical subnets exist, with some logical partitions on each subnet.

### Objective

The objective of this scenario is to configure multiple networks to share a single physical Ethernet adapter. Systems on the same subnet are required to be on the same VLAN and therefore have the same VLAN ID, which allows communication without having to go through the router. The separation in the subnets is achieved by ensuring that the systems on the two subnets have different VLAN IDs.

### Prerequisites and assumptions

- The Hardware Management Console (HMC) was set up. To view the PDF file of the *Installation and Configuration Guide for the Hardware Management Console* (SA76-0084), approximately 3 MB in size, see sa76-0084.pdf .
- You understand the partitioning concepts as described in the *Logical Partitioning Guide*. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf .
- The Virtual I/O Server logical partition has been created and the Virtual I/O Server has been installed. For instructions, see “Installing the Virtual I/O Server and client logical partitions” on page 61.
- You have created the remaining AIX or Linux logical partitions that you want added to the network configuration.
- You have an Ethernet switch and a router ready to add to the configuration.
- You have IP addresses for all logical partitions and systems that will be added to the configuration.

You cannot use VLAN in an Integrated Virtualization Manager environment.

### Configuration steps

The following figure shows the configuration that will be completed during this scenario.

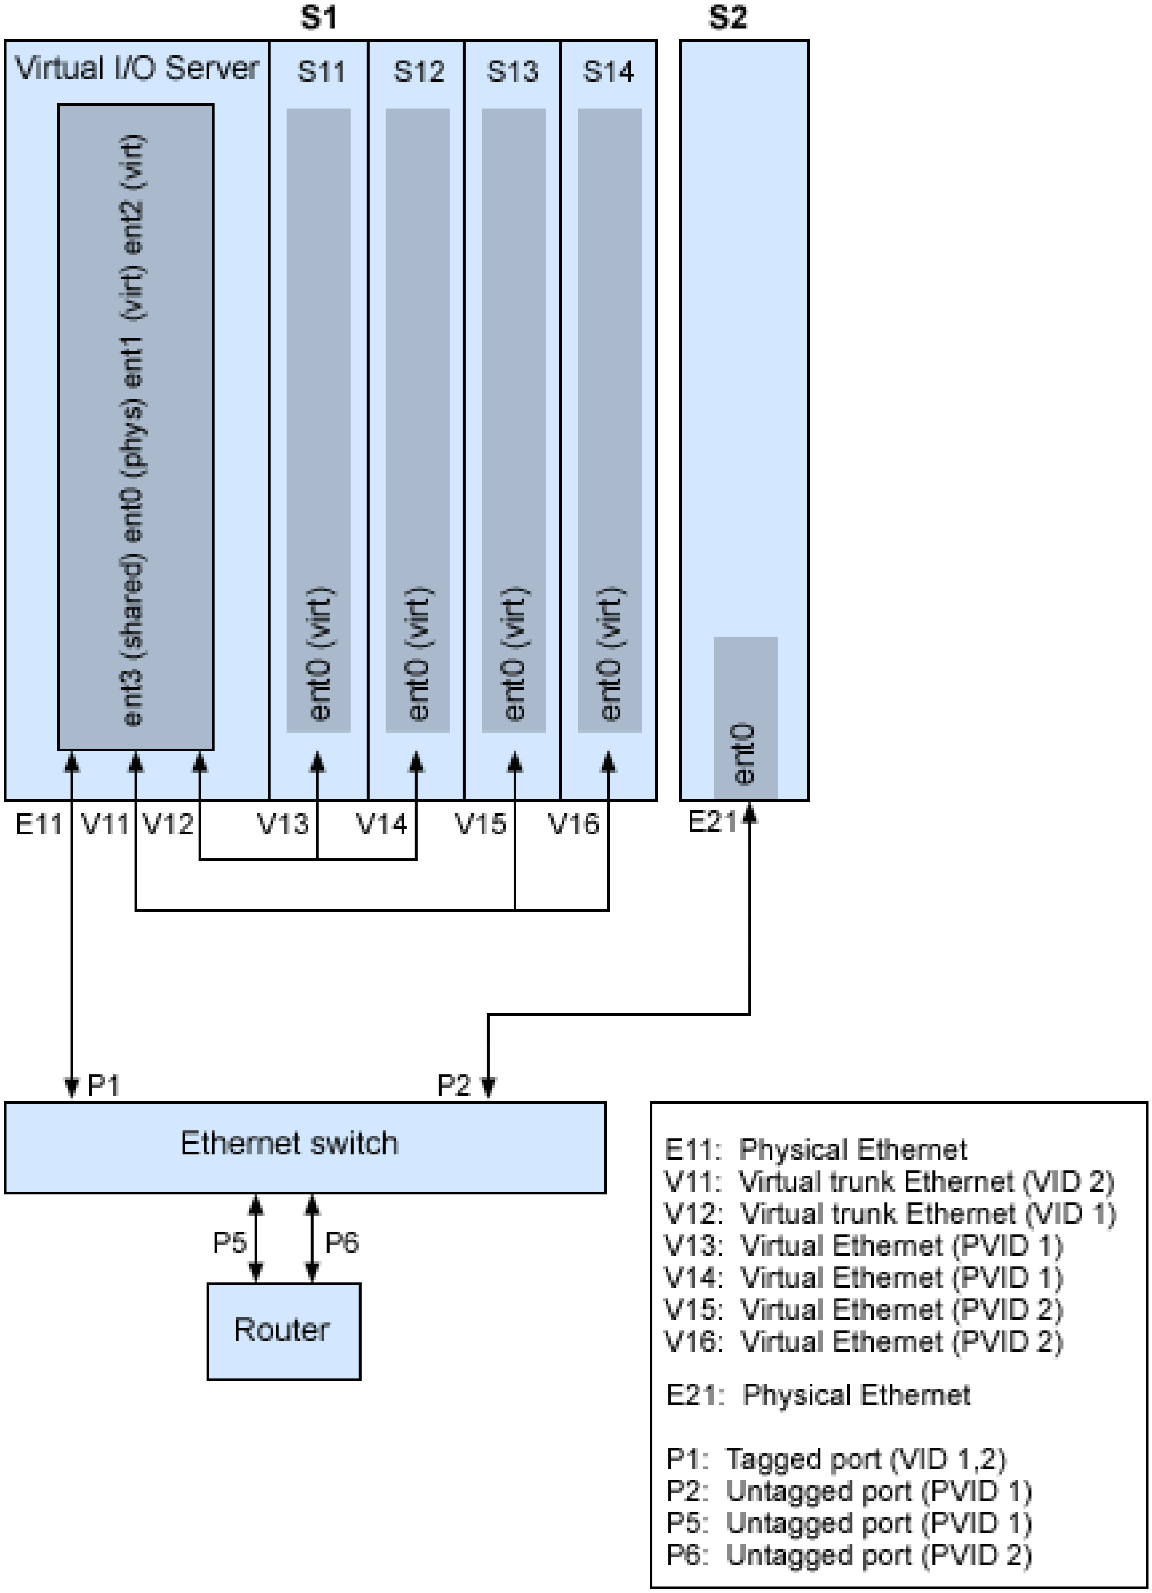
S1
S2
Virtual I/O Server
S11
S12
S13
S14
ent3 (shared) ent0 (phys) ent1 (virt) ent2 (virt)
ent0 (virt)
ent0 (virt)
ent0 (virt)
ent0 (virt)
ent0
E11
V11
V12
V13
V14
V15
V16
E21
P1
P2
Ethernet switch
P5
P6
Router
E11: Physical Ethernet
V11: Virtual trunk Ethernet (VID 2)
V12: Virtual trunk Ethernet (VID 1)
V13: Virtual Ethernet (PVID 1)
V14: Virtual Ethernet (PVID 1)
V15: Virtual Ethernet (PVID 2)
V16: Virtual Ethernet (PVID 2)
E21: Physical Ethernet
P1: Tagged port (VID 1,2)
P2: Untagged port (PVID 1)
P5: Untagged port (PVID 1)
P6: Untagged port (PVID 2)

Using the preceding figure as a guide, follow these steps.

1. Set up the Ethernet switch ports as follows:
   - P1: Tagged port (VID 1, 2)

    - P2: Untagged port (PVID 1)
    - P5: Untagged port (PVID 1)
    - P6: Untagged port (PVID 2)

    For instructions on configuring the ports, see the documentation for your switch.
2. For system S1, use the HMC to create virtual Ethernet adapters for the Virtual I/O Server:
    - Create virtual Ethernet adapter V11 for the Virtual I/O Server with the trunk setting selected and VID set to 2. Specify an unused PVID value. This value is required, even though it will not be used.
    - Create virtual Ethernet adapter V12 for the Virtual I/O Server with the trunk setting selected and VID set to 1. Specify an unused PVID value. This value is required, even though it will not be used.
3. For system S1, use the HMC to create virtual Ethernet adapters for other logical partitions:
    - Create virtual adapters V13 and V14 for logical partitions S11 and S12, respectively, with PVID set to 2 and no additional VIDs.
    - Create virtual adapters V15 and V16 for logical partitions S13 and S14, respectively, with PVID set to 1 and no additional VIDs.
4. For system S1, use the HMC to assign the physical Ethernet adapter (E11) to the Virtual I/O Server and connect the adapter to the Ethernet switch port P1.
5. Using the Virtual I/O Server command-line interface, set up a Shared Ethernet Adapter ent3 with the physical adapter ent0 and virtual adapters ent1 and ent2.
6. Configure IP addresses for the following:
    - S13 (en0), S14 (en0), and S2 (en0) belong to VLAN 1 and are on the same subnet. The router is connected to Ethernet switch port P5.
    - S11 (en0) and S12 (en0) belong to VLAN 2 and are on the same subnet. The router is connected to Ethernet switch port P6.

You can configure the Shared Ethernet Adapter on the Virtual I/O Server logical partition with an IP address. This is required only for network connectivity to the Virtual I/O Server.

As the tagged VLAN network is being used, you must define additional VLAN devices over the Shared Ethernet Adapters before configuring IP addresses.

## Scenario: Configuring Shared Ethernet Adapter failover

Use this article to help you become familiar with typical Shared Ethernet Adapter failover scenario.

**Situation**

You are the system administrator responsible for planning and configuring the network in an environment with the Virtual I/O Server running. You want to provide higher network availability to the client logical partition on the system. This can be accomplished by configuring a backup Shared Ethernet Adapter in a different Virtual I/O Server logical partition.

**Objective**

The objective of this scenario is to configure primary and backup Shared Ethernet Adapters in the Virtual I/O Server logical partitions so that network connectivity in the client logical partitions will not be lost in the case of adapter failure.

**Prerequisites and assumptions**

- The Hardware Management Console (HMC) was set up. To view the PDF file of the *Installation and Configuration Guide for the Hardware Management Console* (SA76-0084), approximately 3 MB in size, see sa76-0084.pdf .

- You understand the partitioning concepts as described in the *Logical Partitioning Guide*. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf .
- Two separate Virtual I/O Server logical partitions have been created and the Virtual I/O Server has been installed in each logical partition. For instructions, see “Installing the Virtual I/O Server and client logical partitions” on page 61.
- You understand what Shared Ethernet Adapter failover is and how it works. See “Shared Ethernet Adapter failover” on page 58.
- You have created the remaining logical partitions that you want added to the network configuration.
- EachVirtual I/O Server logical partition has an available physical Ethernet adapter assigned to it.
- You have IP addresses for all logical partitions and systems that will be added to the configuration.

You cannot use the Integrated Virtualization Manager with multiple Virtual I/O Server logical partitions on the same server.

The following image depicts a configuration where the Shared Ethernet Adapter failover feature is set up. The client logical partitions H1 and H2 are accessing the physical network using the Shared Ethernet Adapters, which are the primary adapters. The virtual Ethernet adapters used in the shared Ethernet setup are configured with the same VLAN membership information (PVID, VID), but have different priorities. A dedicated virtual network forms the control channel and is required to facilitate communication between the primary and backup shared Ethernet device.

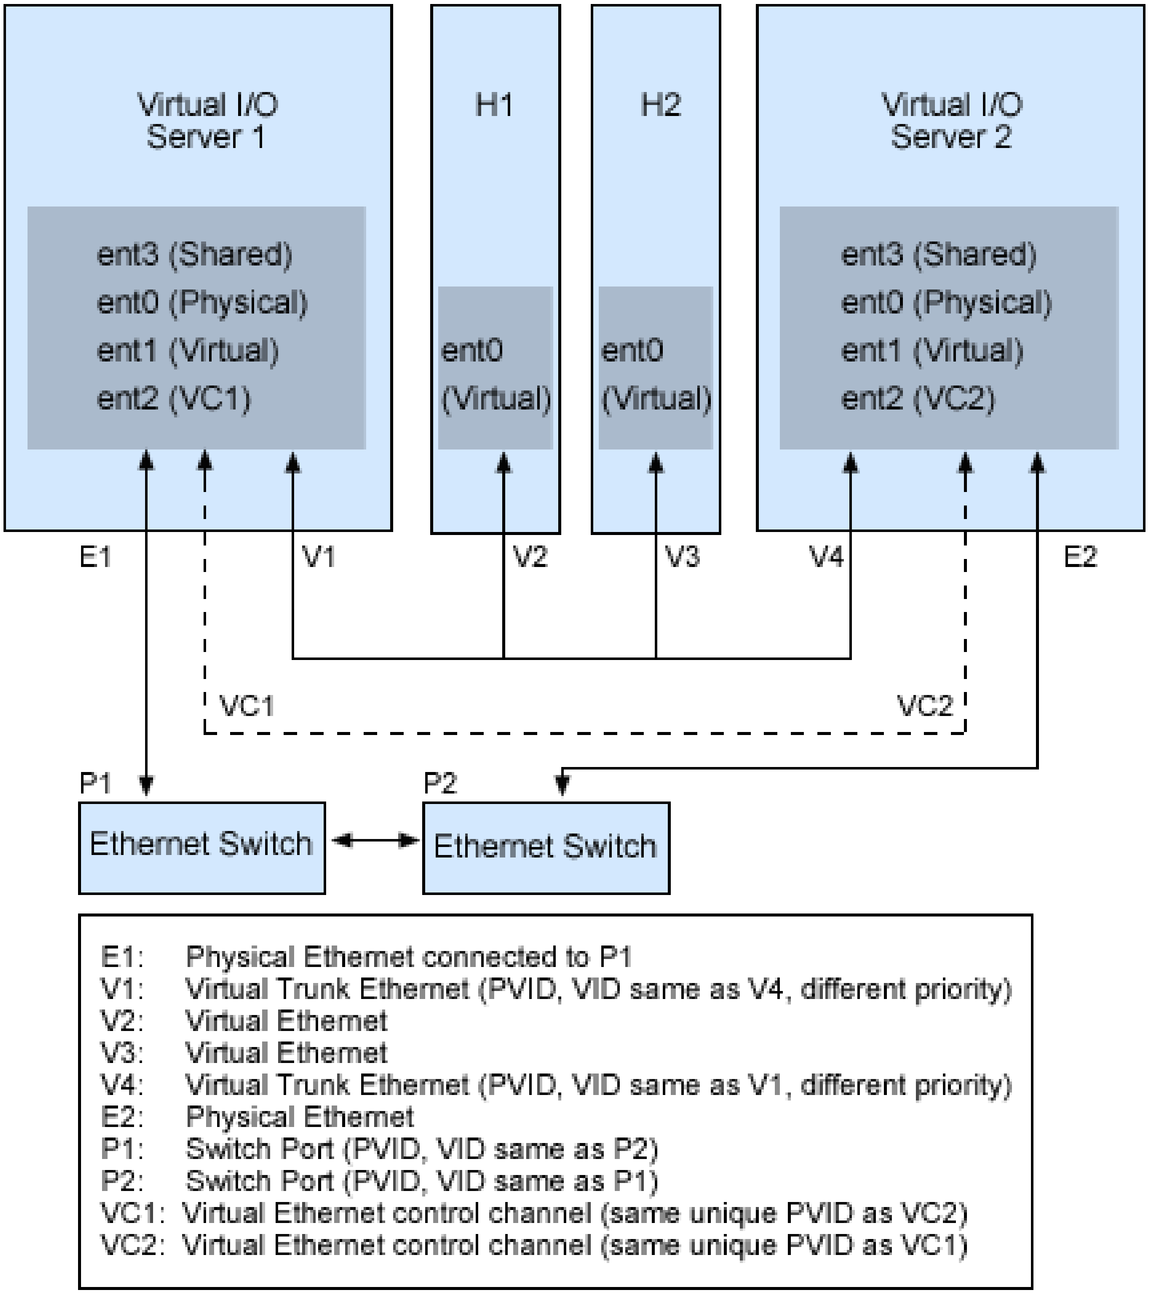
Virtual I/O Server 1
ent3 (Shared)
ent0 (Physical)
ent1 (Virtual)
ent2 (VC1)
H1
ent0 (Virtual)
H2
ent0 (Virtual)
Virtual I/O Server 2
ent3 (Shared)
ent0 (Physical)
ent1 (Virtual)
ent2 (VC2)
E1
V1
V2
V3
V4
E2
VC1
VC2
P1
P2
Ethernet Switch
Ethernet Switch
E1: Physical Ethernet connected to P1
V1: Virtual Trunk Ethernet (PVID, VID same as V4, different priority)
V2: Virtual Ethernet
V3: Virtual Ethernet
V4: Virtual Trunk Ethernet (PVID, VID same as V1, different priority)
E2: Physical Ethernet
P1: Switch Port (PVID, VID same as P2)
P2: Switch Port (PVID, VID same as P1)
VC1: Virtual Ethernet control channel (same unique PVID as VC2)
VC2: Virtual Ethernet control channel (same unique PVID as VC1)

Using the preceding figure as a guide, follow these steps:

1. On the HMC, create the virtual Ethernet adapters following these guidelines:
   - Configure the virtual adapters to be used for data as trunk adapters by selecting the trunk setting.
   - Assign different prioritization values (valid values are 1-15) to each virtual adapter.
   - Configure another virtual Ethernet to be used for the control channel by giving it a unique PVID value. Make sure you use the same PVID when creating this virtual Ethernet for both Virtual I/O Server logical partitions.
2. Using the Virtual I/O Server command line, run the following command to configure the Shared Ethernet Adapter. Run this command on both Virtual I/O Server logical partitions involved in the configuration:

```
mkvdev -sea physical_adapter -vadapter virtual_adapter -default
virtual_adapter\
-defaultid PVID_of_virtual_adapter -attr ha_mode=auto
ctl_chan=control_channel_adapter
```

For example, in this scenario, we ran the following command on both Virtual I/O Server logical partitions:

```
mkvdev -sea ent0 -vadapter ent1 -default ent1 -defaultid 60 -attr ha_mode=auto
ctl_chan=ent2
```

## Scenario: Configuring Network Interface Backup in AIX client logical partitions without VLAN tagging

Use this scenario to become familiar with using a Network Interface Backup configuration in Virtual I/O clients that are running AIX logical partitions and are not configured for VLAN tagging.

**Situation**

In this scenario, you want to configure a highly available virtual environment for your bridged network using the Network Interface Backup (NIB) approach to access external networks from your Virtual I/O clients. You do not plan to use VLAN tagging in your network setup. This approach requires you to configure a second Ethernet adapter on a different VLAN for each client and requires a Link Aggregation adapter with NIB features. This configuration is available for AIX logical partitions.

Typically, a Shared Ethernet Adapter failover configuration is the recommended configuration for most environments because it supports environments with or without VLAN tagging. Also, the NIB configuration is more complex than a Shared Ethernet Adapter failover configuration because it must be implemented on each of the clients. However, Shared Ethernet Adapter failover was not available prior to version 1.2 of Virtual I/O Server, and NIB was the only approach to a highly available virtual environment. Also, you might consider that in an NIB configuration you can distribute clients over both Shared Ethernet Adapters in such a way that half of them will use the first Shared Ethernet Adapter and the other half will use the second Shared Ethernet Adapter as primary adapter.

**Objective**

Create a virtual Ethernet environment using a Network Interface Backup configuration as depicted in the following figure.

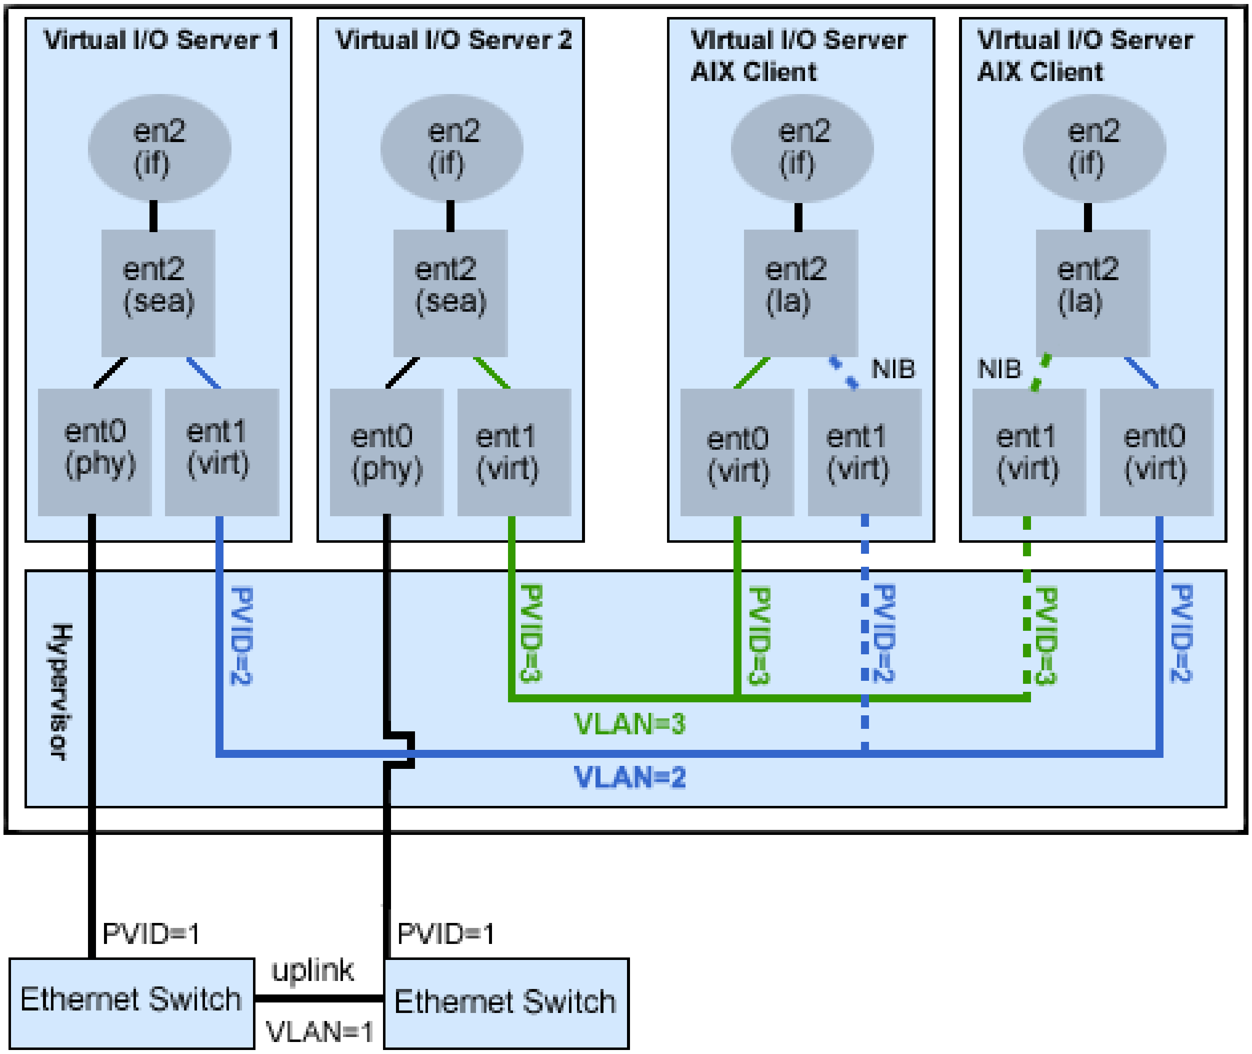

### Prerequisites and assumptions

Before completing the configuration tasks, review the following prerequisites and assumptions.

- The Hardware Management Console (HMC) is already set up. To view the PDF file of the *Installation and Configuration Guide for the Hardware Management Console* (SA76-0084), approximately 3 MB in size, see sa76-0084.pdf.
- Two separate Virtual I/O Server logical partitions have been created and the Virtual I/O Server has been installed in each logical partition. See the instructions in "Installing the Virtual I/O Server and client logical partitions" on page 61.
- You have created the remaining logical partitions that you want added to the network configuration.
- Each Virtual I/O Server logical partition has an available physical Ethernet adapter assigned to it.
- You have IP addresses for all logical partitions and systems that will be added to the configuration.

### Configuration tasks

Using the figure as a guide, complete the following tasks to configure the NIB virtual environment.

1. Create a LAN connection between the Virtual I/O Servers and the external network:

a. Configure a Shared Ethernet Adapter on the primary Virtual I/O Server that bridges traffic between the virtual Ethernet and the external network. See “Configuring a Shared Ethernet Adapter” on page 86.
b. Configure a Shared Ethernet Adapter on the second Virtual I/O Server, as in step 1.

2. For each client logical partition, use the HMC to create a virtual Ethernet whose PVID matches the PVID of the primary Virtual I/O Server. This will be used as the primary adapter.
3. For each client logical partition, use the HMC to create a second virtual Ethernet whose PVID matches the PVID of the second (backup) Virtual I/O Server. This will be used as the backup adapter.
4. Create the Network Interface Backup setup using a Link Aggregation configuration. Make sure that you specify the following items:
   a. Select the primary Ethernet Adapter.
   b. Select the Backup Adapter.
   c. Specify the Internet Address to Ping. Select the IP address or hostname of a host outside of the Virtual I/O Server system that NIB will continuously ping to detect Virtual I/O Server failure.

**Note:** Keep in mind, when you configure NIB with two virtual Ethernet adapters, the internal networks used must stay separated in the hypervisor. You must use different PVIDs for the two adapters in the client and cannot use additional VIDs on them.

## Scenario: Configuring Multi-Path I/O for AIX client logical partitions

Multi-Path I/O (MPIO) helps provide increased availability of virtual SCSI resources by providing redundant paths to the resource. This topic describes how to set up Multi-Path I/O for AIX client logical partitions.

In order to provide MPIO to AIX client logical partitions, you must have two Virtual I/O Server logical partitions configured on your system. This procedure assumes that the disks are already allocated to both the Virtual I/O Server logical partitions involved in this configuration.

To configure MPIO, follow these steps. In this scenario, hdisk5 in the first Virtual I/O Server logical partition, and hdisk7 in the second Virtual I/O Server logical partition, are used in the configuration.

The following figure shows the configuration that will be completed during this scenario.

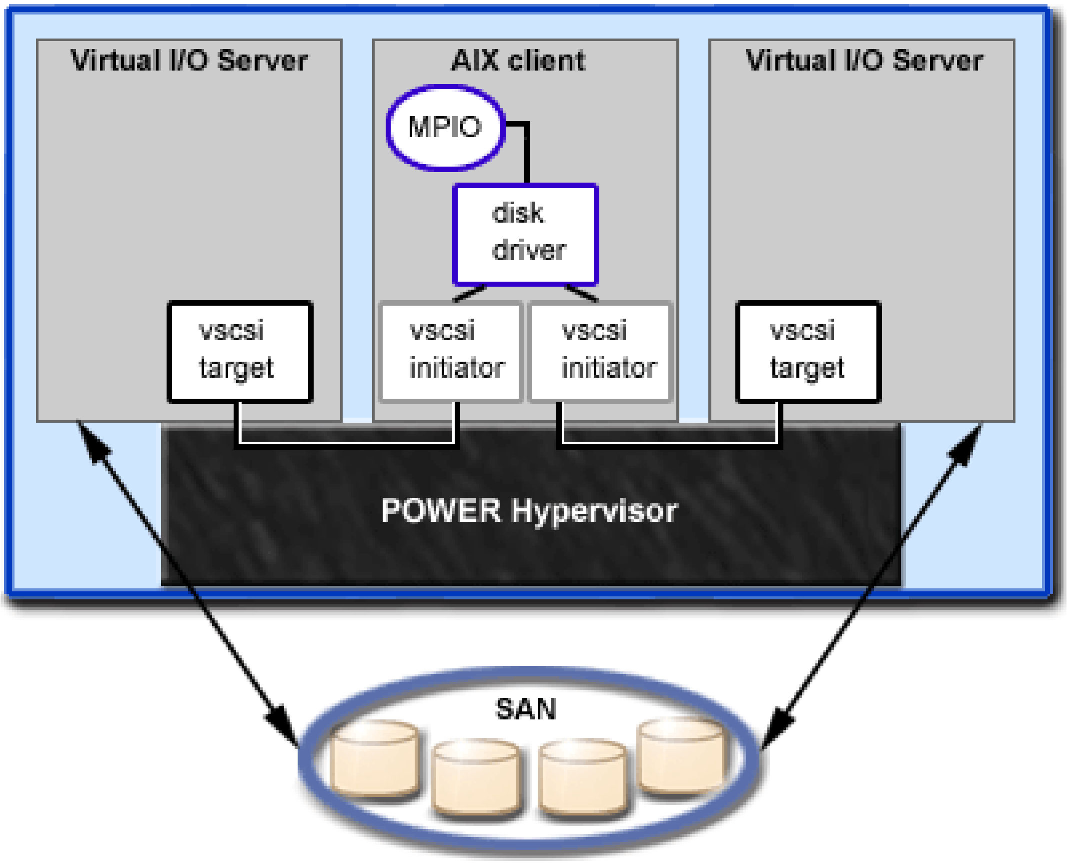

Using the preceding figure as a guide, follow these steps:

1. Using the HMC, create SCSI server adapters on the two Virtual I/O Server logical partitions.
2. Using the HMC, create two virtual client SCSI adapters on the client logical partitions, each mapping to one of the Virtual I/O Server logical partitions.
3. On either of the Virtual I/O Server logical partitions, determine which disks are available by typing `lsdev -type disk`. Your results look similar to the following:

   ```
   name            status     description

   hdisk3          Available  MPIO Other FC SCSI Disk Drive
   hdisk4          Available  MPIO Other FC SCSI Disk Drive
   hdisk5          Available  MPIO Other FC SCSI Disk Drive
   ```

   Select which disk that you want to use in the MPIO configuration. In this scenario, we selected hdisk5.
4. Determine the ID of the disk that you have selected. For instructions, see “Identifying exportable disks” on page 83. In this scenario, the disk does not have an IEEE volume attribute identifier or a unique identifier (UDID), so we determine the physical identifier (PVID) by running the `lspv hdisk5` command. Your results look similar to the following:

   ```
   hdisk5          00c3e35ca560f919                    None
   ```

   The second value is the PVID. In this scenario, the PVID is 00c3e35ca560f919. Note this value.
5. List the attributes of the disk using the **lsdev** command. In this scenario, we typed `lsdev -dev hdisk5 -attr`. Your results look similar to the following

   ```
   ..
   lun_id          0x5463000000000000             Logical Unit Number ID          False
   ..
   ```

```
..
pvid            00c3e35ca560f9190000000000000000 Physical volume identifier      False
..
reserve_policy  single_path                      Reserve Policy                  True
```

Note the values for lun_id and reserve_policy. If the reserve_policy attribute is set to anything other than `no_reserve`, then you must change it. Set the reserve_policy to `no_reserve` by typing `chdev -dev hdiskx -attr reserve_policy=no_reserve`.

6. On the second Virtual I/O Server logical partition, list the physical volumes by typing `lspv`. In the output, locate the disk that has the same PVID as the disk identified previously. In this scenario, the PVID for hdisk7 matched:

   ```
   hdisk7          00c3e35ca560f919                    None
   ```

   **Tip:** Although the PVID values should be identical, the disk numbers on the two Virtual I/O Server logical partitions might vary.

7. Determine if the reserve_policy attribute is set to no_reserve using the **lsdev** command. In this scenario, we typed `lsdev -dev hdisk7 -attr`. You see results similar to the following:

   ```
   ..
   lun_id          0x5463000000000000               Logical Unit Number ID          False
   ..
   pvid            00c3e35ca560f9190000000000000000 Physical volume identifier      False
   ..
   reserve_policy  single_path                      Reserve Policy
   ```

   If the reserve_policy attribute is set to anything other than `no_reserve`, you must change it. Set the reserve_policy to `no_reserve` by typing `chdev -dev hdiskx -attr reserve_policy=no_reserve`.

8. On both Virtual I/O Server logical partitions, use the **mkvdev** to create the virtual devices. In each case, use the appropriate hdisk value. In this scenario, we type the following commands:
   - On the first Virtual I/O Server logical partition, we typed `mkvdev -vdev hdisk5 -vadapter vhost5 -dev vhdisk5`
   - On the second Virtual I/O Server logical partition, we typed `mkvdev -vdev hdisk7 -vadapter vhost7 -dev vhdisk7`

   The same LUN is now exported to the client logical partition from both Virtual I/O Server logical partitions.

9. AIX can now be installed on the client logical partition. For instructions on installing AIX, see Installing AIX in a Partitioned Environment in the servers and AIX Information Center.

10. After you have installed AIX on the client logical partition, check for MPIO by running the following command:

    ```
    lspath
    ```

    You see results similar to the following:

    ```
    Enabled hdisk0 vscsi0
    Enabled hdisk0 vscsi1
    ```

    If one of the Virtual I/O Server logical partitions fails, the results of the lspath command look similar to the following:

    ```
    Failed  hdisk0 vscsi0
    Enabled hdisk0 vscsi1
    ```

    Unless the hcheck_mode and hcheck_interval attributes are set, the state will continue to show `Failed` even after the disk has recovered. To have the state updated automatically, type `chdev -l hdiskx -a hcheck_interval=60 -P`. The client logical partition must be rebooted for this change to take effect.

## Planning for the Virtual I/O Server

Use this topic to help gain an understanding of what to consider when planning for the Virtual I/O Server. In this section, you will find information about planning for the Virtual I/O Server.

### Specifications

This topic defines the range of configuration possibilities, including the minimum number of resources needed and the maximum number of resources allowed.

To activate the Virtual I/O Server, the PowerVM Editions (or Advanced POWER Virtualization) hardware feature is required. A logical partition with enough resources to share with other logical partitions is required. The following is a list of minimum hardware requirements that must be available to create the Virtual I/O Server.

*Table 11. Resources that are required*

| Resource | Requirement |
|---|---|
| Hardware Management Console or Integrated Virtualization Manager | The HMC or Integrated Virtualization Manager is required to create the logical partition and assign resources. |
| Storage adapter | The server logical partition needs at least one storage adapter. |
| Physical disk | The disk must be at least 16 GB. This disk can be shared. |
| Ethernet adapter | If you want to route network traffic from virtual Ethernet adapters to a Shared Ethernet Adapter, you need an Ethernet adapter. |
| Memory | For POWER6 processor-based systems, at least 768 MB of memory is required. For POWER5 processor-based systems, at least 512 MB of memory is required. |
| Processor | At least 0.1 processor is required. |

The following table defines the limitations for storage management.

*Table 12. Limitations for storage management*

| Category | Limit |
|---|---|
| Volume groups | 4096 per system |
| Physical volumes | 1024 per volume group |
| Physical partitions | 1024 per volume group |
| Logical volumes | 1024 per volume group |
| Logical partitions | No limit |

### Limitations and restrictions

Learn about Virtual I/O Server configuration limitations.

Consider the following when implementing virtual SCSI:

- Virtual SCSI supports the following connection standards for backing devices: fibre channel, SCSI, SCSI RAID, iSCSI, SAS, SATA, USB, and IDE.
- The SCSI protocol defines mandatory and optional commands. While virtual SCSI supports all of the mandatory commands, not all of the optional commands are supported.
- There are performance implications when you use virtual SCSI devices. Because the client/server model is made up of layers of function, using virtual SCSI can consume additional processor cycles when processing I/O requests.

- The Virtual I/O Server is a dedicated logical partition, to be used only for Virtual I/O Server operations. Other applications cannot run in the Virtual I/O Server logical partition.
- If there is a resource shortage, performance degradation might occur. If a Virtual I/O Server is serving many resources to other logical partitions, ensure that enough processor power is available. In case of high workload across virtual Ethernet adapters and virtual disks, logical partitions might experience delays in accessing resources.
- Logical volumes and files exported as virtual SCSI disks are always configured as single path devices on the client logical partition.
- Logical volumes or files exported as virtual SCSI disks that are part of the root volume group (rootvg) are not persistent if you reinstall the Virtual I/O Server. However, they are persistent if you update the Virtual I/O Server to a new service pack. Therefore, before reinstalling the Virtual I/O Server, ensure that you back up the corresponding clients' virtual disks. When exporting logical volumes, it is best to export logical volumes from a volume group other than the root volume group. When exporting files, it is best to create file storage pools and the virtual media repository in a parent storage pool other than the root volume group.

Consider the following when implementing virtual adapters:

- Only Ethernet adapters can be shared. Other types of network adapters cannot be shared.
- IP forwarding is not supported on the Virtual I/O Server.
- The maximum number of virtual adapters can be any value from 2 to 65,536. However, if you set the maximum number of virtual adapters to a value higher than 1024, the logical partition might fail to activate or the server firmware might require more system memory to manage the virtual adapters.

The Virtual I/O Server supports client logical partitions running the following operating systems on the following POWER6 processor-based servers.

*Table 13. Operating system support for Virtual I/O Server client logical partitions*

| Operating system | POWER6 processor-based servers |
|---|---|
| AIX 5.3 or later | All POWER6 processor-based servers |
| SUSE Linux Enterprise Server 10 Service Pack 2 or later | • 19F/HA<br>• 25F/2A |
| SUSE Linux Enterprise Server 10 Service Pack 1 | • 03E/4A<br>• 04E/8A<br>• 17M/MA |
| Red Hat Enterprise Linux version 5.2 | • 19F/HA<br>• 25F/2A |
| Red Hat Enterprise Linux version 5.1 | • 03E/4A<br>• 04E/8A<br>• 17M/MA |
| Red Hat Enterprise Linux version 4.7 | 19F/HA |
| Red Hat Enterprise Linux version 4.6 | 25F/2A |
| Red Hat Enterprise Linux version 4.5 | • 03E/4A<br>• 04E/8A<br>• 17M/MA |

The Virtual I/O Server supports client logical partitions running the following operating systems on POWER5 processor-based servers:

- AIX 5.3 (or later)

- SUSE Linux Enterprise Server 9 (or later)
- SUSE Linux Enterprise Server 10 (or later)
- Red Hat Enterprise Linux version 4 (or later)
- Red Hat Enterprise Linux version 5 (or later)

## Capacity planning

This topic includes capacity-planning considerations for the Virtual I/O Server, including information about hardware resources and limitations.

Client logical partitions might use virtual devices, dedicated devices, or a combination of both. Before you begin to configure and install the Virtual I/O Server and client logical partitions, plan what resources each logical partition will use. Throughput requirements and overall workload must be considered when deciding whether to use virtual or dedicated devices and when allocating resources to the Virtual I/O Server. Compared to dedicated SCSI disks, virtual SCSI disks might achieve similar throughput numbers depending on several factors, including workload and virtual SCSI resources. However, virtual SCSI devices generally have higher processor utilization when compared with directly attached storage.

### Planning for virtual SCSI

Find capacity-planning and performance information for virtual SCSI.

Different I/O subsystems have different performance qualities, as does virtual SCSI. This section discusses the performance differences between physical and virtual I/O. The following topics are described in this section:

**Virtual SCSI latency:**

Find information about virtual SCSI latency.

I/O latency is the amount of time that passes between the initiation and completion of a disk I/O operation. For example, consider a program that performs 1000 random disk I/O operations, one at a time. If the time to complete an average operation is 6 milliseconds, the program runs in no fewer than 6 seconds. However, if the average response time is reduced to 3 milliseconds, the run time might be reduced by 3 seconds. Applications that are multithreaded or use asynchronous I/O might be less sensitive to latency, but in most circumstances, lower latency can help improve performance.

Because virtual SCSI is implemented as a client and server model, there is some latency that does not exist with directly attached storage. The latency might range from 0.03 to 0.06 milliseconds per I/O operation depending primarily on the block size of the request. The average latency is comparable for both physical disk and logical volume-backed virtual drives. The latency experienced when using a Virtual I/O Server in a shared-processor logical partition can be higher and more variable than using a Virtual I/O Server in a dedicated logical partition. For additional information about the performance differences between dedicated logical partitions and shared-processor logical partitions, see "Virtual SCSI sizing considerations" on page 48.

The following table identifies latency (in milliseconds) for different block-size transmissions on both physical disk and logical-volume-backed virtual SCSI disks.

*Table 14. Increase in disk I/O response time based on block size (in milliseconds)*

| Backing type | 4 K | 8 K | 32 K | 64 K | 128 K |
|---|---|---|---|---|---|
| Physical disk | 0.032 | 0.033 | 0.033 | 0.040 | 0.061 |
| Logical volume | 0.035 | 0.036 | 0.034 | 0.040 | 0.063 |

The average disk-response time increases as the block size increases. The latency increases for a virtual SCSI operation are relatively greater on smaller block sizes because of their shorter response time.

**Virtual SCSI bandwidth:**

View information about virtual SCSI bandwidth.

I/O bandwidth is the maximum amount of data that can be read or written to a storage device in a unit of time. Bandwidth can be measured from a single thread or from a set of threads running concurrently. Although many customer applications are more sensitive to latency than bandwidth, bandwidth is crucial for many typical operations, such as backing up and restoring persistent data.

The following table compares the results of bandwidth tests for virtual SCSI and physical I/O performance. In the tests, a single thread operates sequentially on a constant file that is 256 MB in size with a Virtual I/O Server running in a dedicated partition. More I/O operations are issued when reading or writing to the file using a small block size as compared to a larger block size. The test was conducted using a storage server with feature code 6239 (type 5704/0625) and a 2-gigabit Fibre Channel adapter attached to one RAID0 LUN that is composed of 5 physical disks from a DS4400 disk system (formerly a FAStT700). The table shows the comparison of measured bandwidth in megabytes per second (MB/s) using virtual SCSI and local attachment for reads with varying block sizes of operations. The difference between virtual I/O and physical I/O in these tests is attributable to the increased latency when using virtual I/O. Because of the larger number of operations, the bandwidth measured with small block sizes is lower than with large block sizes.

*Table 15. Physical and virtual SCSI bandwidth comparison (in MB/s)*

| I/O type | 4 K | 8 K | 32 K | 64 K | 128 K |
|---|---|---|---|---|---|
| Virtual | 20.3 | 35.4 | 82.6 | 106.8 | 124.5 |
| Physical | 24.3 | 41.7 | 90.6 | 114.6 | 132.6 |

**Virtual SCSI sizing considerations:**

Understand the processor and memory-sizing considerations when implementing virtual SCSI .

When you are designing and implementing a virtual SCSI application environment, consider the following sizing issues:

- The amount of memory allocated to the Virtual I/O Server
- The processor entitlement of the Virtual I/O Server
- Whether the Virtual I/O Server is run as a shared-processor logical partition or as a dedicated processor logical partition

The processor impacts of using virtual I/O on the client are insignificant. The processor cycles run on the client to perform a virtual SCSI I/O operation are comparable to that of a locally attached I/O device. Thus, there is no increase or decrease in sizing on the client logical partition for a known task. These sizing techniques do not anticipate combining the function of shared Ethernet with the virtual SCSI server. If the two are combined, consider adding resources to account for the shared Ethernet activity with virtual SCSI .

**Virtual SCSI sizing using dedicated processor logical partitions**

The amount of processor entitlement required for a virtual SCSI server is based on the maximum I/O rates required of it. Because virtual SCSI servers do not normally run at maximum I/O rates all of the time, the use of surplus processor time is potentially wasted when using dedicated processor logical partitions. In the first of the following sizing methodologies, you need a good understanding of the I/O rates and I/O sizes required of the virtual SCSI server. In the second, we will size the virtual SCSI server based on the I/O configuration.

The sizing methodology used is based on the observation that the processor time required to perform an I/O operating on the virtual SCSI server is fairly constant for a given I/O size. It is a simplification to make this statement, because different device drivers have subtly varying efficiencies. However, under most circumstances, the I/O devices supported by the virtual SCSI server are sufficiently similar. The following table shows approximate cycles per second for both physical disk and logical volume operations on a 1.65 Ghz processor. These numbers are measured at the physical processor; simultaneous multithreading (SMT) operation is assumed. For other frequencies, scaling by the ratio of the frequencies (for example, 1.5 Ghz = 1.65 Ghz / 1.5 Ghz × cycles per operation) is sufficiently accurate to produce a reasonable sizing.

*Table 16. Approximate cycles per second on a 1.65 Ghz logical partition*

| Disk type | 4 KB | 8 KB | 32 KB | 64 KB | 128 KB |
|---|---|---|---|---|---|
| Physical disk | 45,000 | 47,000 | 58,000 | 81,000 | 120,000 |
| Logical volume | 49,000 | 51,000 | 59,000 | 74,000 | 105,000 |

Consider a Virtual I/O Server that uses three client logical partitions on physical disk-backed storage. The first client logical partition requires a maximum of 7,000 8-KB operations per second. The second client logical partition requires a maximum of 10,000 8-KB operations per second. The third client logical partition requires a maximum of 5,000 128-KB operations per second. The number of 1.65 Ghz processors for this requirement is approximately ((7,000 × 47,000 + 10,000 × 47,000 + 5,000 × 120,000) / 1,650,000,000) = 0.85 processors, which rounds up to a single processor when using a dedicated processor logical partition.

If the I/O rates of the client logical partitions are not known, you can size the Virtual I/O Server to the maximum I/O rate of the storage subsystem attached. The sizing could be biased toward small I/O operations or large I/O operations. Sizing to maximum capacity for large I/O operations will balance the processor capacity of the Virtual I/O Server to the potential I/O bandwidth of the attached I/O. The negative aspect of this sizing methodology is that, in nearly every case, more processor entitlement will be assigned to the Virtual I/O Server than it will typically consume.

Consider a case in which a Virtual I/O Server manages 32 physical SCSI disks. An upper limit of processors required can be established based on assumptions about the I/O rates that the disks can achieve. If it is known that the workload is dominated by 8096-byte operations that are random, then assume that each disk is capable of approximately 200 disk I/O operations per second (15k rpm drives). At peak, the Virtual I/O Server would need to serve approximately 32 disks × 200 I/O operations per second × 47,000 cycles per operation, resulting in a requirement for approximately 0.19 processor performance. Viewed another way, a Virtual I/O Server running on a single processor should be capable of supporting more than 150 disks doing 8096-byte random I/O operations.

Alternatively, if the Virtual I/O Server is sized for maximum bandwidth, the calculation results in a higher processor requirement. The difference is that maximum bandwidth assumes sequential I/O. Because disks are more efficient when they are performing large, sequential I/O operations than they are when performing small, random I/O operations, a higher number of I/O operations per second can be performed. Assume that the disks are capable of 50 MB per second when doing 128 KB I/O operations. That situation implies each disk could average 390 disk I/O operations per second. Thus, the amount of processing power necessary to support 32 disks, each doing 390 I/O operations per second with an operation cost of 120,000 cycles (32 × 390 × 120,000 / 1,650,000,000) results in approximately 0.91 processors. Consequently, a Virtual I/O Server running on a single processor should be capable of driving approximately 32 fast disks to maximum throughput.

**Virtual SCSI server sizing using shared processor logical partitions**

Defining virtual SCSI servers in shared processor logical partitions allows more specific processor resource sizing and potential recovery of unused processor time by uncapped logical partitions. However, using shared-processor logical partitions for virtual SCSI servers can frequently increase I/O response time and make for somewhat more complex processor entitlement sizings.

The sizing methodology should be based on the same operation costs for dedicated logical partition I/O servers, with added entitlement for running in shared-processor logical partitions. Configure the Virtual I/O Server as uncapped, so that, if the Virtual I/O Server is undersized, there is opportunity to get more processor time to serve I/O operations.

Because I/O latency with virtual SCSI can vary due to a number of conditions, consider the following if a logical partition has high I/O requirements:

- Configure the logical partition with physical I/O if the configuration allows.
- In most cases, the Virtual I/O Server logical partition can use a shared, uncapped processor.

**Virtual SCSI server memory sizing**

Memory sizing in virtual SCSI is simplified because there is no caching of file data in the memory of the virtual SCSI server. Because there is no data caching, the memory requirements for the virtual SCSI server are fairly modest. With large I/O configurations and very high data rates, a 1 GB memory allocation for the virtual SCSI server is likely to be sufficient. For low I/O rate situations with a small number of attached disks, 512 MB will most likely suffice.

## Planning for Shared Ethernet Adapters

Use this section to find capacity-planning and performance information for Shared Ethernet Adapter. This section contains planning information and performance considerations for using Shared Ethernet Adapters on the Virtual I/O Server.

**Network requirements:**

This topic includes information you need in order to accurately size your Shared Ethernet Adapter environment.

To plan for using Shared Ethernet Adapters, you must determine your network needs. This section gives overview information of what should be considered when sizing the Shared Ethernet Adapter environment. Sizing the Virtual I/O Server for the Shared Ethernet Adapter involves the following factors:

- Defining the target bandwidth (MB per second), or transaction rate requirements (operations per second). The target performance of the configuration must be determined from your workload requirements.
- Defining the type of workload (streaming or transaction oriented).
- Identifying the maximum transmission unit (MTU) size that will be used (1500 or jumbo frames).
- Determining if the Shared Ethernet Adapter will run in a threaded or nonthreaded environment.
- Knowing the throughput rates that various Ethernet adapters can provide (see Adapter selection).
- Knowing the processor cycles required per byte of throughput or per transaction (see Processor allocation).

**Bandwidth requirement**

The primary consideration is determining the target bandwidth on the physical Ethernet adapter of the Virtual I/O Server. This will determine the rate that data can be transferred between the Virtual I/O Server and the client logical partitions. After the target rate is known, the correct type and number of network adapters can be selected. For example, Ethernet adapters of various speeds could be used. One or more adapters could be used on individual networks, or they could be combined using Link Aggregation (or EtherChannel).

**Workload type**

The type of workload to be performed must be considered, whether it is streaming of data for workloads such as file transfer, data backup, or small transaction workloads, such as remote procedure calls. The streaming workload consists of large, full-sized network packets and associated small, TCP acknowledgment packets. Transaction workloads typically involve smaller packets or might involve small requests, such as a URL, and a larger response, such as a Web page. A Virtual I/O Server will need to frequently support streaming and small packet I/O during various periods of time. In that case, approach the sizing from both models.

**MTU size**

The MTU size of the network adapters must also be considered. The standard Ethernet MTU is 1500 bytes. Gigabit Ethernet and 10 gigabit Ethernet can support 9000-byte MTU jumbo frames. Jumbo frames might reduce the processor cycles for the streaming types of workloads. However, for small workloads, the larger MTU size might not help reduce processor cycles.

**Threaded or nonthreaded environment**

Use threaded mode when virtual SCSI will be run on the same Virtual I/O Server logical partition as Shared Ethernet Adapter. Threaded mode helps ensure that virtual SCSI and the Shared Ethernet Adapter can share the processor resource appropriately. However, threading increases instruction-path length, which uses additional processor cycles. If the Virtual I/O Server logical partition will be dedicated to running shared Ethernet devices (and associated virtual Ethernet devices) only, the adapters should be configured with threading disabled. For more information, see “Processor allocation” on page 53.

**Adapter throughput**

Knowing the throughput capability of different Ethernet adapters can help you determine which adapters to use as Shared Ethernet Adapters and how many adapters to use. For more information, see “Adapter selection.”

**Processor entitlement**

You must determine how much processor power is required to move data through the adapters at the desired rate. Networking device drivers are typically processor-intensive. Small packets can come in at a faster rate and use more processor cycles than larger packet workloads. Larger packet workloads are typically limited by network wire bandwidth and come in at a slower rate, thus requiring less processor power than small packet workloads for the amount of data transferred.

### Adapter selection:

Use this section to find the attributes and performance characteristics of various types of Ethernet adapters to help you select which adapters to use in your environment.

This section provides approximate throughput rates for various Ethernet adapters set at various MTU sizes. Use this information to determine which adapters will be needed to configure a Virtual I/O Server. To make this determination, you must know the desired throughput rate of the client logical partitions.

Following are general guidelines for network throughput. These numbers are not specific, but they can serve as a general guideline for sizing. In the following tables, the 100 MB, 1 GB, and 10 GB speeds are rounded down for estimating.

*Table 17. Simplex (one direction) streaming rates*

| Adapter speed | Approximate throughput rate |
|---|---|
| 10 Mb Ethernet | 1 MB/second |
| 100 Mb Ethernet | 10 MB/second |
| 1000 Mb Ethernet (GB Ethernet) | 100 MB/second |
| 10000 Mb Ethernet (10 GB Ethernet, Host Ethernet Adapter or Integrated Virtual Ethernet) | 1000 MB/second |

*Table 18. Full duplex (two direction) streaming rates on full duplex network*

| Adapter speed | Approximate throughput rate |
|---|---|
| 10 Mb Ethernet | 2 MB/second |
| 100 Mb Ethernet | 20 MB/second |
| 1000 Mb Ethernet (Gb Ethernet) | 150 MB/second |
| 10000 Mb Ethernet (10 Gb Ethernet, Host Ethernet Adapter or Integrated Virtual Ethernet) | 1500 MB/second |

The following tables list maximum network payload speeds, which are user payload data rates that can be obtained by sockets-based programs for applications that are streaming data. The rates are a result of the network bit rate, MTU size, physical level overhead (such as interframe gaps and preamble bits), data link headers, and TCP/IP headers. A gigahertz-speed processor is assumed. These numbers are optimal for a single LAN. If your network traffic is going through additional network devices, your results might vary.

In the following tables, raw bit rate is the physical media bit rate and does not reflect interframe gaps, preamble bits, data link headers, and trailers. Interframe gaps, preamble bits, data link headers, and trailers can all reduce the effective usable bit rate of the wire.

Single direction (simplex) TCP streaming rates are rates that can be achieved by sending data from one machine to another in a memory-to-memory test. Full-duplex media can usually perform slightly better than half-duplex media because the TCP acknowledgment packets can flow without contending for the same wire that the data packets are flowing on.

*Table 19. Single direction (simplex) TCP streaming rates*

| Network type | Raw bit rate (Mb) | Payload rate (Mb) | Payload rate (MB) |
|---|---|---|---|
| 10 Mb Ethernet, Half Duplex | 10 | 6 | 0.7 |
| 10 Mb Ethernet, Full Duplex | 10 (20 Mb full duplex) | 9.48 | 1.13 |
| 100 Mb Ethernet, Half Duplex | 100 | 62 | 7.3 |
| 100 Mb Ethernet, Full Duplex | 100 (200 Mb full duplex) | 94.8 | 11.3 |
| 1000 Mb Ethernet, Full Duplex, MTU 1500 | 1000 (2000 Mb full duplex) | 948 | 113 |
| 1000 Mb Ethernet, Full Duplex, MTU 9000 | 1000 (2000 Mb full duplex) | 989 | 117.9 |

*Table 19. Single direction (simplex) TCP streaming rates (continued)*

| Network type | Raw bit rate (Mb) | Payload rate (Mb) | Payload rate (MB) |
|---|---|---|---|
| 1000 Mb Ethernet, Full Duplex, Host Ethernet Adapter (or Integrated Virtual Ethernet) MTU 1500 | 10000 | 9479 | 1130 |
| 1000 Mb Ethernet, Full Duplex, Host Ethernet Adapter (or Integrated Virtual Ethernet) MTU 9000 | 10000 | 9899 | 1180 |

Full-duplex TCP streaming workloads have data streaming in both directions. Workloads that can send and receive packets concurrently can take advantage of full duplex media. Some media, for example Ethernet in half-duplex mode, cannot send and receive concurrently, thus they will not perform any better, and can usually degrade performance, when running duplex workloads. Duplex workloads will not increase at a full doubling of the rate of a simplex workload because the TCP acknowledgment packets returning from the receiver must now compete with data packets flowing in the same direction.

*Table 20. Two direction (duplex) TCP streaming rates*

| Network type | Raw bit rate (Mb) | Payload rate (Mb) | Payload rate (MB) |
|---|---|---|---|
| 10 Mb Ethernet, Half Duplex | 10 | 5.8 | 0.7 |
| 10 Mb Ethernet, Full Duplex | 10 (20 Mb full duplex) | 18 | 2.2 |
| 100 Mb Ethernet, Half Duplex | 100 | 58 | 7 |
| 100 Mb Ethernet, Full Duplex | 100 (200 Mb full duplex) | 177 | 21.1 |
| 1000 Mb Ethernet, Full Duplex, MTU 1500 | 1000 (2000 Mb full duplex) | 1470 (1660 peak) | 175 (198 peak) |
| 1000 Mb Ethernet, Full Duplex, MTU 9000 | 1000 (2000 Mb full duplex) | 1680 (1938 peak) | 200 (231 peak) |
| 10000 Mb Ethernet, Host Ethernet Adapter (or Integrated Virtual Ethernet) Full Duplex, MTU 1500 | 10000 | 14680 (15099 peak) | 1750 (1800 peak) |
| 10000 Mb Ethernet, Host Ethernet Adapter (or Integrated Virtual Ethernet) Full Duplex, MTU 9000 | 10000 | 16777 (19293 pack) | 2000 (2300 peak) |

**Note:**

1. Peak numbers represent optimal throughput with multiple TCP sessions running in each direction. Other rates are for a single TCP session.
2. 1000 MB Ethernet (gigabit Ethernet) duplex rates are for the PCI-X adapter in PCI-X slots.
3. Data rates are for TCP/IP using the IPv4 protocol. Adapters with MTU set to 9000 have RFC 1323 enabled.

**Processor allocation:**

This section contains processor-allocation guidelines for both dedicated processor logical partitions and shared processor logical partitions.

Because Ethernet running MTU size of 1500 bytes consumes more processor cycles than Ethernet running Jumbo frames (MTU 9000), the guidelines are different for each situation. In general, the processor utilization for large packet workloads on jumbo frames is approximately half that required for MTU 1500.

If MTU is set to 1500, provide one processor (1.65 Ghz) per Gigabit Ethernet adapter to help reach maximum bandwidth. This equals ten 100-Mb Ethernet adapters if you are using smaller networks. For smaller transaction workloads, plan to use one full processor to drive the Gigabit Ethernet workload to maximum throughput. For example, if two Gigabit Ethernet adapters will be used, allocate up to two processors to the logical partition.

If MTU is set to 9000 (jumbo frames), provide 50% of one processor (1.65 Ghz) per Gigabit Ethernet adapter to reach maximum bandwidth. Small packet workloads should plan to use one full processor to drive the Gigabit Ethernet workload. Jumbo frames have no effect on the small packet workload case.

**Shared Ethernet Adapter using a dedicated processor logical partition**

The sizing provided is divided into two workload types: TCP streaming and TCP request and response. Both MTU 1500 and MTU 9000 networks were used in the sizing, which is provided in terms of machine cycles per byte of throughput for streaming or per transaction for request/response workloads.

The data in the following tables was derived using the following formula:

(number of processors × processor_utilization × processor clock frequency) / Throughput rate in bytes per second or transaction per second = cycles per Byte or transaction.

For the purposes of this test, the numbers were measured on a logical partition with one 1.65 Ghz processor with simultaneous multi-threading (SMT) enabled.

For other processor frequencies, the numbers in these tables can be scaled by the ratio of the processor frequencies for approximate values to be used for sizing. For example, for a 1.5 Ghz processor speed, use 1.65/1.5 × cycles per byte value from the table. This example would result in a value of 1.1 times the value in the table, thus requiring 10% more cycles to adjust for the 10% slower clock rate of the 1.5 Ghz processor.

To use these values, multiply your required throughput rate (in bytes or transactions) by the cycles per byte value in the following tables. This result will give you the required machine cycles for the workload for a 1.65 Ghz speed. Then adjust this value by the ratio of the actual machine speed to this 1.65 Ghz speed. To find the number of processors, divide the result by 1,650,000,000 cycles (or the cycles rate if you adjusted to a different speed machine). You would need the resulting number of processors to drive the workload.

For example, if the Virtual I/O Server must deliver 200 MB of streaming throughput, the following formula would be used:

200 × 1024 × 1024 × 11.2 = 2,348,810,240 cycles / 1,650,000,000 cycles per processor = 1.42 processors.

In round numbers, it would require 1.5 processors in the Virtual I/O Server to handle this workload. Such a workload could then be handled with either a 2-processor dedicated logical partition or a 1.5-processor shared-processor logical partition.

The following tables show the machine cycles per byte for a TCP-streaming workload.

*Table 21. Shared Ethernet with threading option enabled*

| Type of Streaming | MTU 1500 rate and processor utilization | MTU 1500, cycles per byte | MTU 9000 rate and processor utilization | MTU 9000, cycles per byte |
|---|---|---|---|---|
| Simplex | 112.8 MB at 80.6% processor | 11.2 | 117.8 MB at 37.7% processor | 5 |
| Duplex | 162.2 MB at 88.8% processor | 8.6 | 217 MB at 52.5% processor | 3.8 |

*Table 22. Shared Ethernet with threading option disabled*

| Type of Streaming | MTU 1500 rate and processor utilization | MTU 1500, cycles per byte | MTU 9000 rate and processor utilization | MTU 9000, cycles per byte |
|---|---|---|---|---|
| Simplex | 112.8 MB at 66.4% processor | 9.3 | 117.8 MB at 26.7% processor | 3.6 |
| Duplex | 161.6 MB at 76.4% processor | 7.4 | 216.8 MB at 39.6% processor | 2.9 |

The following tables show the machine cycles per transaction for a request and response workload. A transaction is defined as a round-trip request and reply size.

*Table 23. Shared Ethernet with threading option enabled*

| Size of transaction | Transactions per second and Virtual I/O Server utilization | MTU 1500 or 9000, cycles per transaction |
|---|---|---|
| Small packets (64 bytes) | 59,722 TPS at 83.4% processor | 23,022 |
| Large packets (1024 bytes) | 51,956 TPS at 80% processor | 25,406 |

*Table 24. Shared Ethernet with threading option disabled*

| Size of transaction | Transactions per second and Virtual I/O Server utilization | MTU 1500 or 9000, cycles per transaction |
|---|---|---|
| Small packets (64 bytes) | 60,249 TPS at 65.6% processor | 17,956 |
| Large packets (1024 bytes) | 53,104 TPS at 65% processor | 20,196 |

The preceding tables demonstrate that the threading option of the shared Ethernet adds approximately 16 – 20% more machine cycles per transaction for MTU 1500 streaming, and approximately 31% to 38% more machine cycles per transaction for MTU 9000. The threading option adds more machine cycles per transaction at lower workloads due to the threads being started for each packet. At higher workload rates, like full duplex or the request and response workloads, the threads can run longer without waiting and being redispatched. The thread option is a per-shared Ethernet option that can be configured by Virtual I/O Server commands. Disable the thread option if the shared Ethernet is running in a Virtual I/O Server logical partition by itself (without virtual SCSI in the same logical partition).

You can enable or disable threading using the **-attr thread** option of the mkvdev command. To enable threading, use the `-attr thread=1` option. To disable threading, use the `-attr thread=0` option. For example, the following command disables threading for Shared Ethernet Adapter `ent1`:

```
mkvdev -sea ent1 -vadapter ent5 -default ent5 -defaultid 1 -attr thread=0
```

### Sizing a Virtual I/O Server for shared Ethernet on a shared processor logical partition

Creating a shared-processor logical partition for a Virtual I/O Server can be done if the Virtual I/O Server is running slower-speed networks (for example 10/100 Mb) and a full processor logical partition is not needed. It is recommended that this be done only if the Virtual I/O Server workload is less than half a processor or if the workload is inconsistent. Configuring the Virtual I/O Server logical partition as uncapped might also allow it to use more processor cycles as needed to handle inconsistent throughput. For example, if the network is used only when other processors are idle, the Virtual I/O Server logical partition might be able to use other machine cycles and could be created with minimal processor to handle light workload during the day but the uncapped processor could use more machine cycles at night.

If you are creating a Virtual I/O Server in a shared-processor logical partition, add additional processor entitlement as a sizing contingency.

**Memory allocation:**

Find information about memory allocation and sizing.

In general, 512 MB of memory per logical partition is sufficient for most configurations. Enough memory must be allocated for the Virtual I/O Server data structures. Ethernet adapters and virtual devices use dedicated receive buffers. These buffers are used to store the incoming packets, which are then sent over the outgoing device.

A physical Ethernet adapter typically uses 4 MB for MTU 1500 or 16 MB for MTU 9000 for dedicated receive buffers for gigabit Ethernet. Other Ethernet adapters are similar. Virtual Ethernet, typically uses 6 MB for dedicated receive buffers. However, this number can vary based on workload. Each instance of a physical or virtual Ethernet would need memory for this number of buffers. In addition, the system has an mbuf buffer pool per processor that is used if additional buffers are needed. These mbufs typically occupy 40 MB.

## Redundancy considerations

Redundancy options are available at several levels in the virtual I/O environment. Multipathing, mirroring, and RAID redundancy options exist for the Virtual I/O Server and some client logical partitions. Ethernet Link Aggregation (also called EtherChannel) is also an option for the client logical partitions, and the Virtual I/O Server provides Shared Ethernet Adapter failover. There is also support for node failover (HACMP™) for nodes using virtual I/O resources.

This section contains information about redundancy for both the client logical partitions and the Virtual I/O Server. While these configurations help protect from the failure of one of the physical components, such as a disk or network adapter, the might cause the client logical partition to lose access to its devices if the Virtual I/O Server fails. The Virtual I/O Server can be made redundant by running a second instance of it in another logical partition. When running two instances of the Virtual I/O Server, you can use LVM mirroring, multipath I/O, network interface backup, or multipath routing with dead gateway detection in the client logical partition to provide highly available access to virtual resources hosted in separate Virtual I/O Server logical partitions.

### Client logical partitions

This topic includes redundancy considerations for client logical partitions. MPIO, HACMP, and mirroring for the client logical partition are discussed.

**Multipath I/O:**

View Multipath I/O (MPIO) information for client logical partitions.

Multiple virtual SCSI client adapters in a client logical partition can access the same disk through multiple Virtual I/O Server logical partitions. This section describes a virtual SCSI multipath device configuration. If correctly configured, the client recognizes the disk as a multipath device.

Not all virtual SCSI devices are capable of MPIO. To create an MPIO configuration, the exported device at the Virtual I/O Server must conform to the following rules:

- The device must be backed by a physical volume. Logical volume-backed virtual SCSI devices are not supported in an MPIO configuration.
- The device must be accessible from multiple Virtual I/O Server logical partitions.
- The device must be an MPIO-capable device.

  **Note:** MPIO-capable devices are those that contain a unique identifier (UDID) or IEEE volume identifier. For instructions about how to determine whether disks have a UDID or IEEE volume identifier, see "Identifying exportable disks" on page 83.

When setting up an MPIO configuration for virtual SCSI devices on the client logical partition, you must consider the reservation policy of the device on the Virtual I/O Server. To use an MPIO configuration at the client, none of the virtual SCSI devices on the Virtual I/O Server can be reserving the virtual SCSI device. Ensure the **reserve_policy** attribute of the device is set to `no_reserve`. To determine the reserve policy of a device, type the following command:

```
lsdev -dev diskdevicename -attr reserve_policy
```

If the **reserve_policy** value is anything other than `no_reserve`, it must be changed so that you can use the device in an MPIO configuration on the client logical partition. To set the attribute, use the following command:

```
chdev -dev diskdevicename -attr reserve_policy=no_reserve
```

Failover is the only supported behavior for MPIO virtual SCSI disks on the client logical partition.

**Related tasks**

"Scenario: Configuring Multi-Path I/O for AIX client logical partitions" on page 42
Multi-Path I/O (MPIO) helps provide increased availability of virtual SCSI resources by providing redundant paths to the resource. This topic describes how to set up Multi-Path I/O for AIX client logical partitions.

#### Mirroring for client logical partitions:

Achieve mirroring for client logical partitions by using two virtual SCSI adapters.

The client partition can mirror its logical volumes using two virtual SCSI client adapters. Each of these adapters should be assigned to separate Virtual I/O Server partitions. The two physical disks are each attached to a separate Virtual I/O Server partition and made available to the client partition through a virtual SCSI server adapter. This configuration protects virtual disks in a client partition against the failure of any of the following:

- One physical disk
- One physical adapter
- One Virtual I/O Server

The performance of your system might be impacted when using a RAID 1 configuration.

#### High Availability Cluster Multi-Processing:

Learn about High Availability Cluster Multi-Processing (HACMP) in the Virtual I/O Server.

#### HACMP and virtual SCSI

Be aware of the following considerations when implementing HACMP and virtual SCSI:

- The volume group must be defined as Enhanced Concurrent Mode. Enhanced Concurrent Mode is the preferred mode for sharing volume groups in HACMP clusters because volumes are accessible by multiple HACMP nodes. If file systems are used on the standby nodes, those file systems are not mounted until the point of failover. If shared volumes are accessed directly (without file systems) in Enhanced Concurrent Mode, these volumes are accessible from multiple nodes, and as a result, access must be controlled at a higher layer.

- If any one cluster node accesses shared volumes through virtual SCSI, then all nodes must. This means that disks cannot be shared between a logical partition using virtual SCSI and a node directly accessing those disks.
- All volume group configuration and maintenance on these shared disks is done from the HACMP nodes, not from the Virtual I/O Server.

**HACMP and virtual Ethernet**

Be aware of the following considerations when implementing HACMP and virtual Ethernet:

- IP Address Takeover (IPAT) by way of aliasing must be used. IPAT by way of Replacement and MAC Address Takeover are not supported.
- Avoid using the HACMP PCI Hot Plug facility in a Virtual I/O Server environment. PCI Hot Plug operations are available through the Virtual I/O Server. When an HACMP node is using virtual I/O, the HACMP PCI Hot Plug facility is not meaningful because the I/O adapters are virtual rather than physical.
- All virtual Ethernet interfaces defined to HACMP should be treated as single-adapter networks. In particular, you must use the **ping_client_list** attribute to monitor and detect failure of the network interfaces.
- If the Virtual I/O Server has multiple physical interfaces on the same network, or if there are two or more HACMP nodes using the Virtual I/O Server in the same frame, HACMP is not informed of, and does not react to, single physical interface failures. This does not limit the availability of the entire cluster because the Virtual I/O Server routes traffic around the failure.
- If the Virtual I/O Server has only a single physical interface on a network, failure of that physical interface is detected by HACMP. However, that failure isolates the node from the network.

**Link Aggregation or EtherChannel devices:**

A Link Aggregation, or EtherChannel, device is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters can then act as a single Ethernet device. Link Aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

For example, `ent0` and `ent1` can be aggregated to `ent3`. The system considers these aggregated adapters as one adapter, and all adapters in the Link Aggregation device are given the same hardware address, so they are treated by remote systems as if they are one adapter.

Link Aggregation can help provide more redundancy because individual links might fail, and the Link Aggregation device will fail over to another adapter in the device to maintain connectivity. For example, in the previous example, if `ent0` fails, the packets are automatically sent on the next available adapter, `ent1`, without disruption to existing user connections. `ent0` automatically returns to service on the Link Aggregation device when it recovers.

You can configure a Shared Ethernet Adapter to use a Link Aggregation, or EtherChannel, device as the physical adapter.

**Shared Ethernet Adapter failover:**

Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server logical partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption.

A Shared Ethernet Adapter is comprised of a physical adapter (or several physical adapters grouped under a Link Aggregation device) and one or more virtual Ethernet adapters. It can provide layer 2 connectivity to multiple client logical partitions through the virtual Ethernet adapters.

The Shared Ethernet Adapter failover configuration uses the priority value given to the virtual Ethernet adapters during their creation to determine which Shared Ethernet Adapter will serve as the primary and which will serve as the backup. The Shared Ethernet Adapter that has the virtual Ethernet configured with the numerically lower priority value will be used preferentially as the primary adapter. For the purpose of communicating between themselves to determine when a failover should take place, Shared Ethernet Adapters in failover mode use a VLAN dedicated for such traffic, called the *control channel*. For this reason, a virtual Ethernet (created with a PVID that is unique on the system) must be specified as the control channel virtual Ethernet when each Shared Ethernet Adapter is created in failover mode. Using the control channel, the backup Shared Ethernet Adapter is notified when the primary adapter fails, and network traffic from the client logical partitions is sent over the backup adapter. If and when the primary Shared Ethernet Adapter recovers from its failure, it again begins actively bridging all network traffic.

A Shared Ethernet Adapter in failover mode might optionally have more than one trunk virtual Ethernet. In this case, all the virtual Ethernet adapters in a Shared Ethernet Adapter must have the same priority value. Also, the virtual Ethernet adapter used specifically for the control channel does not need to have the trunk adapter setting enabled. The virtual Ethernet adapters used for the control channel on each Shared Ethernet Adapter in failover mode must have an identical PVID value, and that PVID value must be unique in the system, so that no other virtual Ethernet adapters on the same system are using that PVID.

To ensure prompt recovery times, when you enable the Spanning Tree Protocol on the switch ports connected to the physical adapters of the Shared Ethernet Adapter, you can also enable the portfast option on those ports. The portfast option allows the switch to immediately forward packets on the port without first completing the Spanning Tree Protocol. (Spanning Tree Protocol blocks the port completely until it is finished.)

The Shared Ethernet Adapter is designed to prevent network loops. However, as an additional precaution, you can enable Bridge Protocol Data Unit (BPDU) Guard on the switch ports connected to the physical adapters of the Shared Ethernet Adapter. BPDU Guard detects looped Spanning Tree Protocol BPDU packets and shuts down the port. This helps prevent broadcast storms on the network. A *broadcast storm* is a situation where one message that is broadcast across a network results in multiple responses. Each response generates more responses, causing excessive transmission of broadcast messages. Severe broadcast storms can block all other network traffic, but they can usually be prevented by carefully configuring a network to block illegal broadcast messages.

**Note:** When the Shared Ethernet Adapter is using GARP VLAN Registration Protocol (GVRP), it generates BPDU packets, which causes BPDU Guard to shut down the port unnecessarily. Therefore, when the Shared Ethernet Adapter is using GVRP, do not enable BPDU Guard.

For information about how to enable the Spanning Tree Protocol, the portfast option, and BPDU Guard on the ports, see the documentation provided with the switch.

**Related tasks**

"Scenario: Configuring Shared Ethernet Adapter failover" on page 37
Use this article to help you become familiar with typical Shared Ethernet Adapter failover scenario.

## Virtual I/O Server logical partition

Redundancy options for the Virtual I/O Server include multi-pathing, Redundant Array of Independent Disks (RAID) configurations, and Link Aggregation (or EtherChannel).

**Multipathing:**

Multipathing for the physical storage within the Virtual I/O Server provides failover physical path redundancy and load-balancing. The multipathing solutions available in the Virtual I/O Server include MPIO as well as solutions provided by the storage vendors.

**RAID:**

Redundant Array of Independent Disks (RAID) solutions provide for device level redundancy within the Virtual I/O Server. Some RAID options, such as LVM mirroring and striping, are provided by the Virtual I/O Server software, while other RAID options are made available by the physical storage subsystem.

**Link Aggregation or EtherChannel devices:**

A Link Aggregation, or EtherChannel, device is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters can then act as a single Ethernet device. Link Aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

For example, `ent0` and `ent1` can be aggregated to `ent3`. The system considers these aggregated adapters as one adapter, and all adapters in the Link Aggregation device are given the same hardware address, so they are treated by remote systems as if they are one adapter.

Link Aggregation can help provide more redundancy because individual links might fail, and the Link Aggregation device will fail over to another adapter in the device to maintain connectivity. For example, in the previous example, if `ent0` fails, the packets are automatically sent on the next available adapter, `ent1`, without disruption to existing user connections. `ent0` automatically returns to service on the Link Aggregation device when it recovers.

You can configure a Shared Ethernet Adapter to use a Link Aggregation, or EtherChannel, device as the physical adapter.

## Security considerations

Review the security considerations for virtual SCSI, virtual Ethernet, and Shared Ethernet Adapter and the additional security options available.

Systems allow cross-partition device sharing and communication. Functions such as dynamic LPAR, shared processors, virtual networking, virtual storage, and workload management all require facilities to ensure that system-security requirements are met. Cross-partition and virtualization features are designed to not introduce any security exposure beyond what is implied by the function. For example, a virtual LAN connection would have the same security considerations as a physical network connection. Carefully consider how to utilize cross-partition virtualization features in high-security environments. Any visibility between logical partitions must be manually created through administrative system-configuration choices.

Using virtual SCSI, the Virtual I/O Server provides storage to client logical partitions. However, instead of SCSI or fiber cable, the connection for this functionality is done by the firmware. The virtual SCSI device drivers of the Virtual I/O Server and the firmware ensure that only the system administrator of the Virtual I/O Server has control over which logical partitions can access data on Virtual I/O Server storage devices. For example, a client logical partition that has access to a logical volume `lv001` exported by the Virtual I/O Server logical partition cannot access `lv002`, even if it is in the same volume group.

Similar to virtual SCSI, the firmware also provides the connection between logical partitions when using virtual Ethernet. The firmware provides the Ethernet switch functionality. The connection to the external network is provided by the Shared Ethernet Adapter function on the Virtual I/O Server. This part of the Virtual I/O Server acts as a layer-2 bridge to the physical adapters. A VLAN ID tag is inserted into every Ethernet frame. The Ethernet switch restricts the frames to the ports that are authorized to receive frames with that VLAN ID. Every port on an Ethernet switch can be configured to be a member of several

VLANs. Only the network adapters, both virtual and physical, that are connected to a port (virtual or physical) that belongs to the same VLAN can receive the frames. The implementation of this VLAN standard ensures that the logical partitions cannot access restricted data.

## Installing the Virtual I/O Server and client logical partitions

Find instructions for installing the Virtual I/O Server and client logical partitions by deploying a system plan or manually creating the logical partition and logical partition profiles and installing the Virtual I/O Server and client operating systems.

These instructions apply to installing the Virtual I/O Server and client logical partitions on a system that is managed by a Hardware Management Console (HMC). If you plan to install the Virtual I/O Server on a system that is not managed by an HMC, then you need to install the Integrated Virtualization Manager. For instructions, see “Installing the Integrated Virtualization Manager” on page 168.

The installation procedures vary depending on the following factors:

- The version of HMC attached to the managed system on which you plan to install the Virtual I/O Server and client logical partitions. HMC version 7 displays a different interface than prior versions of the HMC. HMC version 7 also provides the ability to deploy a system plan that includes the Virtual I/O Server and client logical partitions.
- Whether you plan to deploy a system plan that includes the Virtual I/O Server and client logical partitions. When you deploy a system plan, the HMC automatically performs the following tasks based on the information provided in the system plan:
  - Creates the Virtual I/O Server logical partition and logical partition profile.
  - Installs the Virtual I/O Server and provisions virtual resources.
  - Creates the client logical partitions and logical partition profiles.
  - Installs the AIX and Linux operating systems on client logical partitions. The HMC must be at V7R3.3.0, or later.

### Installing the Virtual I/O Server and client logical partitions by deploying a system plan

When you deploy a system plan that includes the Virtual I/O Server and, optionally, client logical partitions, the Deploy System Plan wizard creates the Virtual I/O Server logical partition and the logical partition profile, and installs the Virtual I/O Server and client logical partitions.

Before you start, ensure that you meet the following requirements:

- The system to which you plan to deploy the system plan is managed by a Hardware Management Console (HMC).
- The HMC is at version 7 or later. If the HMC is at a version 6 or earlier, then you cannot deploy a system plan. You must manually create the Virtual I/O Server logical partition and logical partition profile and install the Virtual I/O Server. For instructions, see “Installing the Virtual I/O Server manually using the HMC version 6” on page 72.
- If you plan to deploy different entities of the Virtual I/O Server configuration at different times, ensure that the HMC is at version V7R3.3.0, or later. (Virtual I/O Server entities include Shared Ethernet Adapters, EtherChannel adapters, or Link Aggregation devices, storage pools, and backing devices.) If the HMC is not at V7R3.3.0, or later, system plans that include the Virtual I/O Server can be deployed only to new systems, or to systems that do not already have a Virtual I/O Server logical partition configured. (The Virtual I/O Server can be installed, but not configured.) More specifically, no Virtual I/O Server entities can be configured on the managed system, including Shared Ethernet Adapters, EtherChannel adapters, or Link Aggregation devices, storage pools, and backing devices.
- If you plan to deploy a system plan that includes AIX or Linux installation information for at least one client logical partition, ensure that you meet the following requirements:

  - The HMC must be at V7R3.3.0.
  - The client logical partition does not have an operating system already installed. The HMC installs AIX and Linux on client logical partitions that do not already have an operating system installed. If the client logical partition already has an operating system installed, the HMC does not deploy the operating system specified in the system plan.

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

Use these instructions to enter the PowerVM Editions (or Advanced POWER Virtualization) activation code using the Hardware Management Console (HMC) version 7, or later.

If PowerVM Editions is not enabled on your system, you can use the HMC to enter the activation code that you received when you ordered the feature.

Use the following procedure to enter the activation code for the PowerVM Standard Edition and the PowerVM Enterprise Edition. For information about the PowerVM Editions, see “PowerVM Editions” on page 2.

To enter your activation code, follow these steps:

1. In the Navigation area, expand **Systems Management**.
2. Select **Servers**.
3. In the contents area, select the managed system on which you plan to use PowerVM Editions. For example, this might be the system on which you plan to install the Virtual I/O Server, or it might be the system in which you plan to use the Micro-Partitioning technology.
4. Click **Tasks** and select **Capacity on Demand (CoD)** → **Advanced POWER Virtualization** → **Enter Activation Code**.
5. Enter your activation code and click **OK**.

## Importing a system plan into an HMC Version 7

You can import a system-plan file into a Hardware Management Console (HMC) from various types of media, a remote FTP site, or the computer from which you remotely access the HMC. You can then deploy the imported system plan to a system that the HMC manages.

You can import a system-plan file into the HMC from any of the following locations:

- From the computer on which you remotely access the HMC.
- From various media that is mounted on the HMC, such as optical discs or USB drives.
- From a remote site by using FTP. To use this option, you must fulfill the following requirements:
  - The HMC must have a network connection to the remote site.
  - An FTP server must be active on the remote site.
  - Port 21 must be open on the remote site.

**Note:** You cannot import a system plan that has an identical name to any system plan that is available on the HMC.

To import a system-plan file, you must be a super administrator. For more information about user roles, refer to Tasks and roles in the *Operations Guide for the Hardware Management Console and Managed Systems*.

To import a system-plan file into Version 7 of the HMC, complete the following steps:

1. In the navigation area of the HMC, select **System Plans**.
2. In the tasks area, select **Import System Plan**. The Import System Plan window opens.
3. Select the source of the system-plan file that you want to import. Use the following table to complete the appropriate steps for importing the system plan from the selected source location of the file.

| Source of the system plan to import | Complete the following steps: |
|---|---|
| **This computer** | 1. Select **Import from this computer to the HMC**.<br>2. Click **Import** to display the Upload File window.<br>3. Click **Browse**.<br>4. Select the system-plan file that you want to import and click **Open**.<br>5. Click **OK** to upload the file. |
| **Media** | 1. Select **Import from media**.<br>2. In the **System plan file name** field, enter the name of the system-plan file. **Note:** The name of the system-plan file must end with the .sysplan file name suffix and can use alphanumeric characters only.<br>3. In the **Sub-directory on media** field, enter the path in which the system-plan file is located on the media. **Note:** Specify the subdirectory location only, rather than the fully qualified path and file name.<br>4. Click **Import** to display the Select Media Device window.<br>5. Select the media that contains the system-plan file that you want to import.<br>6. Click **OK**. |
| **Remote FTP site** | 1. Select **Import from a remote FTP site**.<br>2. In the **System plan file name** field, enter the name of the system-plan file. **Note:** The name of the system-plan file must end with the .sysplan file name suffix and can use alphanumeric characters only.<br>3. In the **Remote site hostname** field, enter the host name or IP address of the remote FTP site.<br>4. In the **User ID** field, enter the user ID to use to access the remote FTP site.<br>5. In the **Password** field, enter the password to use to access the remote FTP site.<br>6. In the **Remote directory** field, enter the path in which the system-plan file is located on the remote FTP site. If you do not enter a path, the HMC uses the default path specified on the remote FTP site. |

4. Click **Import**. If the HMC returns an error, return to the **Import System Plan** window and verify that the information you entered is correct. If necessary, click **Cancel**, return to step 2, and redo the procedure, ensuring that the information you specify at each step is correct.

**Note:** As an alternative to the HMC Web user interface, you can use the `cpysysplan` command from the HMC command line interface to import a system plan.

When you complete the process of importing the system-plan file, you can deploy the system plan in the system-plan file to a system that the HMC manages. For instructions, see Deploying a system plan by using HMC version 7. If you imported the system-plan file from media, you can unmount the media by using the umount command from the HMC command line interface.

**Related tasks**

"Deploying a system plan by using HMC Version 7"
You can use the Hardware Management Console (HMC) to deploy all or part of a system plan to a managed system.

**Related information**

Operations Guide for the Hardware Management Console and its Managed Systems
This publication provides system administrators and system operators with information about using the Hardware Management Console.

## Deploying a system plan by using HMC Version 7

You can use the Hardware Management Console (HMC) to deploy all or part of a system plan to a managed system.

When you deploy a system plan, the HMC creates logical partitions on the managed system according to the specifications in the system plan. Depending on the contents of the system plan, you can also install operating environments on the logical partitions in the plan and, if the plan contains Virtual I/O Server provisioning information for a logical partition, such as storage assignments, the HMC can make these resource assignments for the logical partition.

**Requirements for deploying a system plan**

Before you deploy a system plan, complete the following tasks:

1. Ensure that the system-plan file exists on the HMC. If the system-plan file does not exist on the HMC, you must import the system-plan file into the HMC. For instructions, see Importing a system plan using HMC version 7.
2. Verify that the physical hardware and any expansion units are connected and are reporting to the server. Each server comes with one logical partition and one partition profile. All of the physical hardware resources on the system are assigned automatically to this logical partition so that you can power on the server and verify that the physical hardware is connected and reporting to the server.
3. Locate the physical disk I/O adapters that belong to each logical partition and verify that the disk drives that are attached to these physical I/O adapters support your desired configuration for each logical partition. The Deploy System Plan Wizard validates only that the physical disk I/O adapters match the system plan. It does not validate that the disk drives are configured for the physical disk I/O adapters. If you are deploying a system plan that you created by using the HMC, verify that the hardware and cabling on the target system is identical to that on the source system.
4. Delete the logical partition that was provided with your server, and delete any other logical partition that is not in the system plan. For instructions, see Deleting a logical partition in the *Logical Partitioning Guide*. The name of the logical partition that was provided with the server is the serial number of the managed system, and the name of the partition profile is *default_profile*.
5. If the system plan includes a Storage Area Network (SAN) or Fibre Channel adapters, ensure that the adapters are cabled and that the SAN is configured.
6. If you plan to deploy the Virtual I/O Server, AIX, or Linux operating environment for a logical partition, ensure that the appropriate installation image is either on the HMC or available to the HMC through a connection to a Network Installation Management (NIM) server. To see the installation images on the HMC, either enter the `OS_install -l` command on the HMC command line or, for HMC V7R3.3.0, use the **Manage Install Resources** task in the HMC Web interface. If the Virtual I/O Server, AIX, or Linux installation image that you need is not listed, complete the following steps to copy an installation image to the HMC:
   a. Obtain a copy of the Virtual I/O Server, AIX, or Linux on DVD or CD, whichever is appropriate for the operating environment. You can use the original installation media or you can contact your marketing representative to obtain another copy. If you cannot obtain a copy of the Virtual I/O Server, AIX, or Linux operating environment, you can deploy the remainder of the system plan and install the Virtual I/O Server, AIX, or Linux operating environment at a later time.

b. Copy the installation files that you need either to a NIM Server or to the HMC itself. To copy the necessary installation files to the HMC, insert the DVD into the DVD drive on the HMC. If the installation files are on CDs, insert the initial CD into the CD drive on the HMC.

   **Note:** You can use the OS_install command to perform this task. From the HMC command line, use the OS_install command to copy the operating environment installation files. For example, you can enter the following command to copy the Virtual I/O Server installation image from the DVD to the HMC:

   ```
   OS_install -o define_resource -a type=VIOS -a version=x.x -a location=/extra/csminstall/vios
   -a source=/dev/cdrom VIOS
   ```

   where *x.x* can be any of the following values: 1.4.1.0, 1.5, or 1.5.2.

c. On the HMC, select **HMC Management** → **Manage Install Resources**, and click **Add**.

d. In the **Add Install Resource** window, select **Create local install resource**, select the appropriate operating environment name and version, and click **OK** to copy the installation files from the installation media to the hard drive of the HMC.

7. If you plan to deploy a system plan that includes the installation of an operating environment for a logical partition, ensure that the **Power off the system after all the logical partitions are powered off** attribute for the managed system is not selected. If this attribute is selected, system plan deployment will fail because the deployment process starts partitions and then powers off partitions as part of installing operating environments. Consequently, the managed system will power off during deployment when the deployment process powers off the partitions. To verify this system attribute, complete these steps:

   a. In the HMC navigation area, select **Systems Management** → **Servers**.

   b. In the Tasks area, click **Properties**. The Properties window for the selected managed system opens.

   c. On the **General** tab, verify that the **Power off the system after all the logical partitions are powered off** attribute is not selected, and click **OK**.

8. For Virtual I/O Server logical partitions previously deployed, ensure that they are active, and that there is a Resource Monitoring and Control (RMC) connection between the HMC and each Virtual I/O Server logical partition. An RMC connection also is necessary to ensure that the HMC can verify the operating environment status for a logical partition. To install an operating environment as part of deploying a system plan, the Deploy System Plan Wizard must be able to determine if the affected logical partition already has an operating environment installed. The wizard can obtain this information from the partition properties on the HMC when a logical partition has been started at least once with an active RMC connection. This ensures that the operating environment status for the logical partition is known to the HMC and ensures that the wizard can determine whether it is appropriate to install the operating environment on the logical partition as specified in the system plan.

9. Ensure that you are not using this HMC or any other HMC that is attached to the managed system to perform any other operations on the managed system.

10. Ensure that you are a super administrator. For information about user roles, refer to Tasks and roles in the *Operations Guide for the Hardware Management Console and Managed Systems*.

11. Ensure that a system plan with Virtual I/O Server partitions is suitable for deployment on the managed system. On HMCs prior to V7R3.3.0, the managed system must not have any Virtual I/O Server entities configured on it for any Virtual I/O Server partitions that exist on it. Specifically, the managed system cannot have any Virtual I/O Server entities configured on it, including shared Ethernet adapters, EtherChannel adapters, or link aggregation devices, storage pools, and backing devices. If you try to deploy a system plan on a system that does not meet all of these requirements, the Deploy System Plan Wizard fails the validation step. On HMC V7R3.3.0, partition validation is more flexible. Consequently, you can deploy a system plan with Virtual I/O Server partitions even if the managed system has Virtual I/O Server partitions with Virtual I/O Server entities configured for them.

### Deploying a system plan

To use the HMC to deploy a system plan on a managed system, complete the following steps:

1. In the navigation area of the HMC, select **System Plans**.
2. In the contents area, select the system plan that you want to deploy.
3. Select **Tasks** → **Deploy system plan**. The Deploy System Plan Wizard starts.
4. On the Welcome page, complete the following steps:
   a. Select the system-plan file that contains the system plan that you want to deploy.
   b. Choose the managed system to which you want to deploy the system plan and click **Next**. If the system plan does not match the managed system to which you want to deploy the plan, the wizard displays a window that informs you of this. Click **OK** to continue or **Cancel** to select a different system plan.

      **Note:** If the system-plan file contains multiple system plans, the wizard provides a step so that you can select a specific system plan from the file. This step does not occur unless there is more than one system plan in the specified file.
5. On the Validation page, complete the following steps:
   a. Wait for the wizard to validate the managed system and its hardware against the system plan. The validation process can take several minutes.
   b. If the validation process completes successfully, click **Next**.
   c. If the validation process fails, correct the problems that the error messages describe, click **Cancel** to exit the wizard, and restart this procedure from the beginning. To help you correct any validation problems, you might want to create a system plan that is based on the current configuration of the managed system. Such a system plan allows you to compare the system plan that you want to deploy with the current configuration of the managed system. You can do this by using the Create System Plan task in the HMC, or you can run the following command from the HMC command line:

      ```
      mksysplan -m name_of_managed_system -f name_of_new_system_plan.sysplan
      ```

      This action creates a new system plan that you can view and compare to the old system plan to help diagnose any problems.

      **Important:** The mksysplan command performs hardware discovery by default on models where the hardware discovery process is available. If you run the hardware discovery process while logical partitions are active, any data about hardware that those active partitions own is removed from the inventory cache. It is recommended that you keep the inventory cache current. Therefore, either use the *-nohwdisc* parameter to have the command run without performing a new hardware discovery, or move all partitions to the inactive state before running this mksysplan command.
6. Optional: On the Partition Deployment page, if you do not want to create all of the logical partitions, partition profiles, virtual adapter types, or virtual adapters in the system plan, clear the boxes in the **Deploy** column beside the logical partitions, partition profiles, virtual adapter types, or virtual adapters that you do not want to create. Virtual serial adapters are required in virtual slots 0 and 1 for each logical partition. You cannot create the logical partition unless you create these virtual serial adapters.
7. Optional: On the Operating Environment Install page, if there is operating environment installation information specified in the system plan, complete the following steps:
   a. Select the operating environments that you want to deploy to the managed system for each logical partition. For HMC V7R3.2.0 or V7R3.1.0, you can deploy only the Virtual I/O Server operating environment. For HMC V7R3.3.0, or later, versions, you also can select to deploy the AIX or Linux operating environments if the system plan contains installation information for them.
   b. Enter the location of the Virtual I/O Server installation image.

c. Enter or change late-binding installation settings for the specified Virtual I/O Server, AIX, or Linux operating environment. Late-binding installation settings are settings that are specific to the installation instance and must be supplied during the installation step to ensure that the settings are accurate for the installation instance. For example, you can enter the IP address of the target logical partition on which you are installing the operating environment.

   **Note:** If you need to use automatic installation files to deploy an operating environment, you cannot add them during the HMC deployment process.

d. Save any changes that you make to late-binding installation settings. You can save them to the current system-plan file or to a new system-plan file.

8. On the Summary page, review the system deployment step order and click **Finish**. The HMC uses the system plan to create the specified logical partitions and to install any specified operating environments. This process can take several minutes.

After you finish the deployment of the system plan, install operating environments and software on the logical partitions, if they did not install as part of system plan deployment.

**Related tasks**

"Importing a system plan into an HMC Version 7" on page 62
You can import a system-plan file into a Hardware Management Console (HMC) from various types of media, a remote FTP site, or the computer from which you remotely access the HMC. You can then deploy the imported system plan to a system that the HMC manages.

**Related information**

Logical Partitioning Guide

Operations Guide for the Hardware Management Console and its Managed Systems
This publication provides system administrators and system operators with information about using the Hardware Management Console.

## Finishing the Virtual I/O Server installation

After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

This procedure assumes that Virtual I/O Server is installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 61.

To finish the installation, complete the following steps:

1. Accept the software maintenance terms and conditions, and the Virtual I/O Server product license. For instructions, see "Viewing and accepting the Virtual I/O Server license" on page 72.
2. Set up remote connections to the Virtual I/O Server. For instructions, see "Connecting to the Virtual I/O Server using OpenSSH" on page 119.
3. Optional: Create the following additional user IDs. After the installation, the only active user ID is the prime administrator (padmin). You can create the following additional user IDs: system administrator, service representative, and development engineer. For information about creating user IDs, see "Managing users on the Virtual I/O Server" on page 124.
4. Configure the TCP/IP connection for the Virtual I/O Server using the mktcpip command. You must complete this task before you can perform any dynamic logical partitioning operations. Alternatively, you can use the configuration assistance menu to configure TCP/IP connections. You can access the configuration assistance menu by running the cfgassist command.

When you are finished, do one of the following tasks:

- If you installed the Virtual I/O Server, client logical partitions, and operating systems by completely deploying a system plan, your setup is complete. For information about how to manage the Virtual I/O Server, see "Managing the Virtual I/O Server" on page 94.

- If you installed the Virtual I/O Server manually using HMC version 6 or version 7, you need to configure the Virtual I/O Server, create client logical partitions, and install client operating systems. For information, see “Configuring the Virtual I/O Server” on page 76 and the *Logical Partitioning Guide*. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf .

## Installing the Virtual I/O Server manually using the HMC version 7

You can create the Virtual I/O Server logical partition and logical partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 7 or later.

Before you start, ensure that the following statements are true:

- The system on which you plan install the Virtual I/O Server is managed by a Hardware Management Console (HMC).
- The HMC is at version 7 or later. If the HMC is at a version 6 or earlier, then see Installing the Virtual I/O Server manually using the HMC version 6.

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

Use these instructions to enter the PowerVM Editions (or Advanced POWER Virtualization) activation code using the Hardware Management Console (HMC) version 7, or later.

If PowerVM Editions is not enabled on your system, you can use the HMC to enter the activation code that you received when you ordered the feature.

Use the following procedure to enter the activation code for the PowerVM Standard Edition and the PowerVM Enterprise Edition. For information about the PowerVM Editions, see “PowerVM Editions” on page 2.

To enter your activation code, follow these steps:

1. In the Navigation area, expand **Systems Management**.
2. Select **Servers**.
3. In the contents area, select the managed system on which you plan to use PowerVM Editions. For example, this might be the system on which you plan to install the Virtual I/O Server, or it might be the system in which you plan to use the Micro-Partitioning technology.
4. Click **Tasks** and select **Capacity on Demand (CoD)** → **Advanced POWER Virtualization** → **Enter Activation Code**.
5. Enter your activation code and click **OK**.

### Creating the Virtual I/O Server logical partition and partition profile using HMC version 7

You can use the Hardware Management Console (HMC) version 7 to create a logical partition and partition profile for the Virtual I/O Server.

Before you start, ensure that the following statements are true:

- You are a super administrator or an operator.
- The PowerVM Editions (or Advanced POWER Virtualization) feature is activated. For instructions, see “Entering the activation code for PowerVM Editions using the HMC version 7” on page 10.

The Virtual I/O Server requires a minimum of 16 GB of disk space.

To create a logical partition and a partition profile on your server using the HMC, follow these steps:

1. In the Navigation area, expand **Systems Management**.
2. Select **Servers**.

3. In the contents area, select the server on which you want to create the partition profile.
4. Click **Tasks** and select **Configuration** → **Create Logical Partition** → **VIO Server**.
5. On the Create Partition page, enter a name and ID for the Virtual I/O Server partition.
6. On the Partition Profile page, complete the following steps:
   a. Enter a profile name for the Virtual I/O Server partition.
   b. Make sure that the **Use all the resources in the system** check box is cleared (not checked).
7. On the Processors page, decide if you want to use shared or dedicated processors (based on your environment) by making the appropriate selection.
8. On the Processing Settings page, enter the appropriate amount of processing units and virtual processors that you want to assign to the Virtual I/O Server partition.
9. On the Memory page, select the appropriate amount of memory that you want to assign to the Virtual I/O Server partition. The required minimum is 512 MB.
10. On the I/O page, select the physical I/O resources that you want in the Virtual I/O Server partition.
11. On the Virtual Adapters page, create the appropriate adapters for your environment.
12. On the Logical Host Ethernet Adapter (LHEA) page, configure one or more LHEAs for the Virtual I/O Server partition. (Host Ethernet Adapter is sometimes referred to as Integrated Virtual Ethernet.)
13. On the Optional Settings page, complete the following steps:
   a. Decide if you want connection monitoring by making the appropriate selection.
   b. If you want the Virtual I/O Server to start when the managed system starts, select the **Automatically start with managed system** option.
   c. Decide if you want to enable redundant error path reporting by making the appropriate selection.
   d. Select the boot mode for the Virtual I/O Server partition. In most cases, the **Normal** boot mode is the appropriate selection.
14. Verify your selections in the Profile Summary window and click **Finish**.

After you create the partition and partition profile, you are ready to install the Virtual I/O Server. For instructions, see one of the following procedures:
- "Installing the Virtual I/O Server from the HMC"
- "Installing the Virtual I/O Server from CD or DVD" on page 70

## Installing the Virtual I/O Server from the HMC

Find instructions for installing the Virtual I/O Server from the HMC by using the installios command.

Before you start, complete the following tasks:
1. Ensure that the following statements are true:
   - There is an HMC attached to the managed system.
   - The Virtual I/O Server logical partition and logical partition profile are created. For instructions, see one of the following tasks:
     - "Creating the Virtual I/O Server logical partition and partition profile using HMC version 7" on page 68
     - "Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6" on page 73
   - The Virtual I/O Server logical partition has at least one Ethernet adapter and a 16 GB disk assigned to it.
   - You have **hmcsuperadmin** authority.
2. Gather the following information:
   - Static IP address for the Virtual I/O Server
   - Subnet mask for the Virtual I/O Server
   - Default gateway for the Virtual I/O Server

To install the Virtual I/O Server, follow these steps:

1. Insert the Virtual I/O Server CD or DVD into the HMC.
2. If you are installing the Virtual I/O Server through the public network interface, continue to step 3. If you are installing the Virtual I/O Server through a private network interface, type the following from the HMC command line:

   ```
   export INSTALLIOS_PRIVATE_IF=interface
   ```

   where *interface* is the network interface through which the installation should take place.
3. From the HMC command line, type:

   ```
   installios
   ```
4. Follow the installation instructions according to the system prompts.

After you install the Virtual I/O Server, finish the installation by checking for updates, setting up remote connections, creating additional user IDs, and so on. For instructions, see “Finishing the Virtual I/O Server installation” on page 67.

## Installing the Virtual I/O Server from CD or DVD

Find instructions for installing the Virtual I/O Server from a CD or DVD device that is attached to the Virtual I/O Server logical partition.

Before you start, ensure that the following statements are true:

- There is an HMC attached to the managed system.
- The Virtual I/O Server logical partition and logical partition profile are created. For instructions, see one of the following tasks:
  - “Creating the Virtual I/O Server logical partition and partition profile using HMC version 7” on page 68
  - “Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6” on page 73
- A CD or DVD optical device is assigned to the Virtual I/O Server logical partition.

To install the Virtual I/O Server from CD or DVD, follow these steps:

1. Activate the Virtual I/O Server logical partition using the HMC version 7 (or later) or HMC version 6 (or earlier):
   - Activate the Virtual I/O Server using the HMC version 7 or later:
     a. Insert the Virtual I/O Server CD or DVD into the Virtual I/O Server logical partition.
     b. In the HMC navigation area, expand **Systems Management** → **Servers**.
     c. Select the server on which the Virtual I/O Server logical partition is located.
     d. In the contents area, select the Virtual I/O Server logical partition.
     e. Click **Tasks** → **Operations** → **Activate**. The Activate Partition menu opens with a selection of logical partition profiles. Ensure the correct profile is highlighted.
     f. Select **Open a terminal window or console session** to open a virtual terminal (vterm) window.
     g. Click **(Advanced)** to open the advanced options menu.
     h. For the boot mode, select **SMS**.
     i. Click **OK** to close the advanced options menu.
     j. Click **OK**. A virtual terminal window opens for the logical partition.
   - Activate the Virtual I/O Server using the HMC version 6 or earlier:
     a. Insert the Virtual I/O Server CD or DVD into the Virtual I/O Server logical partition.
     b. On the HMC, right-click the logical partition to open the menu.
     c. Click **Activate**. The Activate Partition menu opens with a selection of logical partition profiles. Ensure the correct profile is highlighted.

    d. Select **Open a terminal window or console session** to open a virtual terminal (vterm) window.
    e. Click **(Advanced)** to open the advanced options menu.
    f. For the boot mode, select **SMS**.
    g. Click **OK** to close the advanced options menu.
    h. Click **OK**. A virtual terminal window opens for the logical partition.
2. Select the boot device:
    a. Select **Select Boot Options** and press Enter.
    b. Select **Select Install/Boot Device** and press Enter.
    c. Select **Select 1st Boot Device** and press Enter.
    d. Select **CD/DVD** and press Enter.
    e. Select the media type that corresponds to the optical device and press Enter.
    f. Select the device number that corresponds to the optical device and press Enter.
    g. Set the boot sequence to configure the first boot device. The optical device is now the first device in the Current Boot Sequence list.
    h. Exit the SMS menu by pressing the x key, and confirm that you want to exit SMS.
3. Install the Virtual I/O Server:
    a. Select the desired console and press Enter.
    b. Select a language for the BOS menus and press Enter.
    c. Select **Start Install Now with Default Settings** and press Enter.
    d. Select **Continue with Install**. The system will reboot after the installation is complete.

After you install the Virtual I/O Server, finish the installation by checking for updates, setting up remote connects, creating additional user IDs, and so on. For instructions, see "Finishing the Virtual I/O Server installation" on page 67.

## Finishing the Virtual I/O Server installation

After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

This procedure assumes that Virtual I/O Server is installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 61.

To finish the installation, complete the following steps:

1. Accept the software maintenance terms and conditions, and the Virtual I/O Server product license. For instructions, see "Viewing and accepting the Virtual I/O Server license" on page 72.
2. Set up remote connections to the Virtual I/O Server. For instructions, see "Connecting to the Virtual I/O Server using OpenSSH" on page 119.
3. Optional: Create the following additional user IDs. After the installation, the only active user ID is the prime administrator (padmin). You can create the following additional user IDs: system administrator, service representative, and development engineer. For information about creating user IDs, see "Managing users on the Virtual I/O Server" on page 124.
4. Configure the TCP/IP connection for the Virtual I/O Server using the mktcpip command. You must complete this task before you can perform any dynamic logical partitioning operations. Alternatively, you can use the configuration assistance menu to configure TCP/IP connections. You can access the configuration assistance menu by running the cfgassist command.

When you are finished, do one of the following tasks:

- If you installed the Virtual I/O Server, client logical partitions, and operating systems by completely deploying a system plan, your setup is complete. For information about how to manage the Virtual I/O Server, see "Managing the Virtual I/O Server" on page 94.

- If you installed the Virtual I/O Server manually using HMC version 6 or version 7, you need to configure the Virtual I/O Server, create client logical partitions, and install client operating systems. For information, see "Configuring the Virtual I/O Server" on page 76 and the *Logical Partitioning Guide*. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf .

**Viewing and accepting the Virtual I/O Server license:**

You must view and accept the license before using the Virtual I/O Server.

Before you start, ensure that the Virtual I/O Server logical partition profile is created and the Virtual I/O Server is installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 61.

To view and accept the Virtual I/O Server license, complete the following steps:

1. Log in to the Virtual I/O Server using the **padmin** user ID.
2. Choose a new password. The software maintenance terms and conditions appear.
3. View and accept the Virtual I/O Server product license.

   **Note:** If you installed the Virtual I/O Server by deploying a system plan, then you have already accepted the Virtual I/O Server product license and do not need to complete this step.

   a. To view the Virtual I/O Server product license, type `license -ls` on the command line. By default, the license is displayed in English. To change the language in which the license is displayed, follow these steps:

      1) View the list of available locales to display the license by typing the following command:

      ```
      license -ls
      ```

      2) View the license in another language by typing the following command:

      ```
      license -view -lang Name
      ```

      For example, to view the license in Japanese, type the following command:

      ```
      license -view -lang ja_JP
      ```

   b. To accept the Virtual I/O Server product license, type `license -accept` on the command line.
4. In the installation program, English is the default language. If you need to change the language setting for the system, follow these steps:

   a. View the available languages by typing the following command:

   ```
   chlang -ls
   ```

   b. Change the language by typing the following command, replacing *Name* with the name of the language you are switching to:

   ```
   chlang -lang Name
   ```

   **Note:** If the language fileset is not installed, use the **-dev** ***Media*** flag to install it.
   For example, to install and change the language to Japanese, type the following command:

   ```
   chlang -lang ja_JP -dev /dev/cd0
   ```

## Installing the Virtual I/O Server manually using the HMC version 6

You can create the Virtual I/O Server logical partition and logical partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 6 or earlier.

Before you start, ensure that the following statements are true:

- The system on which you plan install the Virtual I/O Server is managed by a Hardware Management Console (HMC).

- The HMC is at version 6 or earlier. If the HMC is at version 7 or later, then see one of the following procedures:
  - "Installing the Virtual I/O Server and client logical partitions by deploying a system plan" on page 61
  - "Installing the Virtual I/O Server manually using the HMC version 7" on page 68

## Entering the activation code for PowerVM Editions using the HMC version 6

The PowerVM Editions (or Advanced POWER Virtualization) activation code is required to install and configure the Virtual I/O Server. You can enter the code using the Hardware Management Console (HMC).

If the PowerVM Editions feature is not enabled on your system, you must use the HMC to enter the activation code that you received when you ordered the feature. This activation code also enables the Micro-Partitioning technology on the system.

To enter your activation code, follow these steps:

1. From the HMC, select the managed system.
2. Select **Manage On Demand Activations**.
3. Select **Virtualization Engine Technologies**.
4. Select **Enter Activation Code**. Type your activation code.

## Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6

You can use the Hardware Management Console (HMC) version 6 to create a logical partition and partition profile for the Virtual I/O Server.

Before you start, ensure that the following statements are true:

- You are a super administrator or an operator.
- The PowerVM Editions (or Advanced POWER Virtualization) feature is activated. For instructions, see "Entering the activation code for PowerVM Editions using the HMC version 6" on page 11.

The Virtual I/O Server requires a minimum of 16 GB of disk space and 512 MB of memory.

To create a logical partition and a partition profile on your server using the HMC, follow these steps:

1. In the Navigation Area, open **Server and Partition**.
2. Select **Server Management**.
3. In the contents area, open the server on which you want to create the partition profile.
4. Right-click **Partitions** and select **Create** → **Logical Partition**.
5. Enter a name for the Virtual I/O Server partition.
6. Select the Virtual I/O Server as the Partition Environment.
7. Based on your environment, decide whether the Virtual I/O Server will be part of a workload management group.
8. Enter a profile name for the Virtual I/O Server partition.
9. Make sure that the **Use all the resources in the system** check box is cleared (not checked).
10. Select the appropriate amount of memory that you want to assign to the Virtual I/O Server partition. The required minimum is 512 MB.
11. Based on your environment, decide if you want to use shared or dedicated processors by making the appropriate selection.
12. Select the physical I/O resources that you want in the Virtual I/O Server partition.
13. Based on your environment, decide if the Virtual I/O Server will use I/O pools by making the appropriate selection.

14. In the Virtual I/O Adapters window, select Yes to indicate that you want to specify virtual adapters.
15. In the Create Virtual I/O Adapters window, create the appropriate adapters for your environment.
16. Based on your environment, decide if you want to specify a power-controlling partition for the Virtual I/O Server partition.
17. Decide if you want connection monitoring by making the appropriate selection.
18. If you want the Virtual I/O Server to start when the managed system starts, select the **Automatically start with managed system** option.
19. Select the boot mode for the Virtual I/O Server partition. In most cases, the **Normal Boot Mode** is the appropriate selection.
20. Verify your selections in the Profile Summary window and click **Finish**.

After creating your logical partition and partition profile, you must install the Virtual I/O Server. For instructions, see one of the following procedures:
- "Installing the Virtual I/O Server from the HMC" on page 69
- "Installing the Virtual I/O Server from CD or DVD" on page 70

## Installing the Virtual I/O Server from the HMC

Find instructions for installing the Virtual I/O Server from the HMC by using the installios command.

Before you start, complete the following tasks:
1. Ensure that the following statements are true:
   - There is an HMC attached to the managed system.
   - The Virtual I/O Server logical partition and logical partition profile are created. For instructions, see one of the following tasks:
     - "Creating the Virtual I/O Server logical partition and partition profile using HMC version 7" on page 68
     - "Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6" on page 73
   - The Virtual I/O Server logical partition has at least one Ethernet adapter and a 16 GB disk assigned to it.
   - You have **hmcsuperadmin** authority.
2. Gather the following information:
   - Static IP address for the Virtual I/O Server
   - Subnet mask for the Virtual I/O Server
   - Default gateway for the Virtual I/O Server

To install the Virtual I/O Server, follow these steps:
1. Insert the Virtual I/O Server CD or DVD into the HMC.
2. If you are installing the Virtual I/O Server through the public network interface, continue to step 3. If you are installing the Virtual I/O Server through a private network interface, type the following from the HMC command line:
   ```
   export INSTALLIOS_PRIVATE_IF=interface
   ```
   where *interface* is the network interface through which the installation should take place.
3. From the HMC command line, type:
   ```
   installios
   ```
4. Follow the installation instructions according to the system prompts.

After you install the Virtual I/O Server, finish the installation by checking for updates, setting up remote connections, creating additional user IDs, and so on. For instructions, see "Finishing the Virtual I/O Server installation" on page 67.

## Installing the Virtual I/O Server from CD or DVD

Find instructions for installing the Virtual I/O Server from a CD or DVD device that is attached to the Virtual I/O Server logical partition.

Before you start, ensure that the following statements are true:

- There is an HMC attached to the managed system.
- The Virtual I/O Server logical partition and logical partition profile are created. For instructions, see one of the following tasks:
  - "Creating the Virtual I/O Server logical partition and partition profile using HMC version 7" on page 68
  - "Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6" on page 73
- A CD or DVD optical device is assigned to the Virtual I/O Server logical partition.

To install the Virtual I/O Server from CD or DVD, follow these steps:

1. Activate the Virtual I/O Server logical partition using the HMC version 7 (or later) or HMC version 6 (or earlier):
   - Activate the Virtual I/O Server using the HMC version 7 or later:
     a. Insert the Virtual I/O Server CD or DVD into the Virtual I/O Server logical partition.
     b. In the HMC navigation area, expand **Systems Management** → **Servers**.
     c. Select the server on which the Virtual I/O Server logical partition is located.
     d. In the contents area, select the Virtual I/O Server logical partition.
     e. Click **Tasks** → **Operations** → **Activate**. The Activate Partition menu opens with a selection of logical partition profiles. Ensure the correct profile is highlighted.
     f. Select **Open a terminal window or console session** to open a virtual terminal (vterm) window.
     g. Click **(Advanced)** to open the advanced options menu.
     h. For the boot mode, select **SMS**.
     i. Click **OK** to close the advanced options menu.
     j. Click **OK**. A virtual terminal window opens for the logical partition.
   - Activate the Virtual I/O Server using the HMC version 6 or earlier:
     a. Insert the Virtual I/O Server CD or DVD into the Virtual I/O Server logical partition.
     b. On the HMC, right-click the logical partition to open the menu.
     c. Click **Activate**. The Activate Partition menu opens with a selection of logical partition profiles. Ensure the correct profile is highlighted.
     d. Select **Open a terminal window or console session** to open a virtual terminal (vterm) window.
     e. Click **(Advanced)** to open the advanced options menu.
     f. For the boot mode, select **SMS**.
     g. Click **OK** to close the advanced options menu.
     h. Click **OK**. A virtual terminal window opens for the logical partition.
2. Select the boot device:
   a. Select **Select Boot Options** and press Enter.
   b. Select **Select Install/Boot Device** and press Enter.
   c. Select **Select 1st Boot Device** and press Enter.
   d. Select **CD/DVD** and press Enter.
   e. Select the media type that corresponds to the optical device and press Enter.
   f. Select the device number that corresponds to the optical device and press Enter.

g. Set the boot sequence to configure the first boot device. The optical device is now the first device in the Current Boot Sequence list.

h. Exit the SMS menu by pressing the x key, and confirm that you want to exit SMS.

3. Install the Virtual I/O Server:
   a. Select the desired console and press Enter.
   b. Select a language for the BOS menus and press Enter.
   c. Select **Start Install Now with Default Settings** and press Enter.
   d. Select **Continue with Install**. The system will reboot after the installation is complete.

After you install the Virtual I/O Server, finish the installation by checking for updates, setting up remote connects, creating additional user IDs, and so on. For instructions, see "Finishing the Virtual I/O Server installation" on page 67.

### Finishing the Virtual I/O Server installation

After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

This procedure assumes that Virtual I/O Server is installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 61.

To finish the installation, complete the following steps:

1. Accept the software maintenance terms and conditions, and the Virtual I/O Server product license. For instructions, see "Viewing and accepting the Virtual I/O Server license" on page 72.
2. Set up remote connections to the Virtual I/O Server. For instructions, see "Connecting to the Virtual I/O Server using OpenSSH" on page 119.
3. Optional: Create the following additional user IDs. After the installation, the only active user ID is the prime administrator (padmin). You can create the following additional user IDs: system administrator, service representative, and development engineer. For information about creating user IDs, see "Managing users on the Virtual I/O Server" on page 124.
4. Configure the TCP/IP connection for the Virtual I/O Server using the mktcpip command. You must complete this task before you can perform any dynamic logical partitioning operations. Alternatively, you can use the configuration assistance menu to configure TCP/IP connections. You can access the configuration assistance menu by running the cfgassist command.

When you are finished, do one of the following tasks:

- If you installed the Virtual I/O Server, client logical partitions, and operating systems by completely deploying a system plan, your setup is complete. For information about how to manage the Virtual I/O Server, see "Managing the Virtual I/O Server" on page 94.
- If you installed the Virtual I/O Server manually using HMC version 6 or version 7, you need to configure the Virtual I/O Server, create client logical partitions, and install client operating systems. For information, see "Configuring the Virtual I/O Server" and the *Logical Partitioning Guide*. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf .

## Configuring the Virtual I/O Server

You need to configure virtual SCSI and virtual Ethernet devices on the Virtual I/O Server. Optionally, you can also configure Tivoli agents and clients and configure the Virtual I/O Server as an LDAP client.

## Configuring virtual SCSI on the Virtual I/O Server

You can configure virtual SCSI devices by deploying a system plan, creating volume groups and logical volumes, and configuring the Virtual I/O Server to support SCSI-2 reserve functions.

Provisioning virtual disk resources occurs on the Virtual I/O Server. Physical disks owned by the Virtual I/O Server can either be exported and assigned to a client logical partition as a whole or can be partitioned into parts, such as logical volumes or files. These logical volumes and files can be exported as virtual disks to one or more client logical partitions. Therefore, by using virtual SCSI, you can share adapters as well as disk devices.

To make a physical volume, logical volume, or file available to a client logical partition requires that it be assigned to a virtual SCSI server adapter on the Virtual I/O Server. The SCSI client adapter is linked to a particular virtual SCSI server adapter in the Virtual I/O Server logical partition. The client logical partition accesses its assigned disks through the virtual SCSI client adapter. The Virtual I/O Server client adapter sees standard SCSI devices and LUNs through this virtual adapter. Assigning disk resources to a SCSI server adapter in the Virtual I/O Server effectively allocates resources to a SCSI client adapter in the client logical partition.

## Creating the virtual target device on the Virtual I/O Server

Creating the virtual target device on the Virtual I/O Server maps the virtual SCSI adapter with the file, logical volume, or physical disk.

With the Virtual I/O Server version 1.5 and later, you can export the following types of physical disks:
- Virtual SCSI disk backed by a physical volume
- Virtual SCSI disk backed by a logical volume
- Virtual SCSI disk backed by a file

After a virtual disk is assigned to a client partition, the Virtual I/O Server must be available before the client logical partitions can access it.

### Creating a virtual target device on the Virtual I/O Server that maps to a physical or logical volume:

You can create a virtual target device on the Virtual I/O Server that maps the virtual SCSI adapter to a physical disk or a logical volume that is based on a volume group.

The following procedure can be repeated to provide additional virtual disk storage to any client logical partition.

Before you start, ensure the following statements are true:
1. At least one physical or logical volume is defined on the Virtual I/O Server. For information, see "Logical volumes" on page 18.
2. The virtual adapters for the Virtual I/O Server and the client logical partitions are created. This usually occurs during the creation of the logical partition profile. For information about creating the logical partition, see Installing the Virtual I/O Server.

To create a virtual target device that maps a virtual SCSI server adapter to a physical or logical volume, complete the following steps:
1. Use the lsdev command to ensure that the virtual SCSI adapter is available. For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   ent3     Available  Virtual I/O Ethernet Adapter (l-lan)
   vhost0   Available  Virtual SCSI Server Adapter
   vhost1   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   vtscsi0  Available  Virtual Target Device - Logical Volume
   vtscsi1  Available  Virtual Target Device - File-backed Disk
   vtscsi2  Available  Virtual Target Device - File-backed Disk
   ```
2. To create a virtual target device, which maps the virtual SCSI server adapter to a physical or logical volume, run the mkvdev command:

```
mkvdev -vdev TargetDevice -vadapter VirtualSCSIServerAdapter
```

Where:

- *TargetDevice* is the name of the target device, as follows:
  - To map a logical volume to the virtual SCSI server adapter, use the name of the logical volume. For example, lv_4G.
  - To map a physical volume to the virtual SCSI server adapter, use `hdiskx`. For example, hdisk5.
  - To map an optical device to the virtual SCSI server adapter, use `cdx`. For example, cd0.
- *VirtualSCSIServerAdapter* is the name of the virtual SCSI server adapter.

The storage is available to the client logical partition either the next time it starts, or the next time the appropriate virtual SCSI client adapter is probed (on a Linux logical partition), or configured (on an AIX logical partition).

3. View the newly created virtual target device by running the lsdev command. For example, running `lsdev -virtual` returns results similar to the following:

```
name     status     description
vhost3   Available  Virtual SCSI Server Adapter
vsa0     Available  LPAR Virtual Serial Adapter
vtscsi0  Available  Virtual Target Device - Logical Volume
```

4. View the logical connection between the newly created devices by running the lsmap command. For example, running `lsmap -vadapter vhost3` returns results similar to the following:

```
SVSA      Physloc                    Client PartitionID
-------------------------------------------------------
vhost3    U9111.520.10DDEEC-V1-C20   0x00000000

VTD                    vtscsi0
Status                 Available
LUN                    0x8100000000000000
Backing device         lv_4G
Physloc
```

The physical location is a combination of the slot number, in this case 20, and the logical partition ID. The storage is now available to the client logical partition either the next time it starts, or the next time the appropriate virtual SCSI client adapter is probed, or configured.

If you later need to remove the virtual target device, you can do so by using the rmvdev command.

**Creating a virtual target device on the Virtual I/O Server that maps to a file or logical volume:**

You can create a virtual target device on the Virtual I/O Server that maps the virtual SCSI adapter to a file or a logical volume that is based on a storage pool.

The following procedure can be repeated to provide additional virtual disk storage to any client logical partition.

Before you start, ensure the following statements are true:

- The Virtual I/O Server is at version 1.5 or later.
- At least one file is defined in a file storage pool, or at least one logical volume is defined in a logical volume storage pool on the Virtual I/O Server. For information, see “Virtual storage” on page 22 and “Storage pools” on page 21.
- The virtual adapters for the Virtual I/O Server and the client logical partitions are created. This usually occurs during the creation of the logical partition profile. For information about creating the logical partition, see Installing the Virtual I/O Server.

To create a virtual target device that maps a virtual SCSI server adapter to a file or logical volume, complete the following steps:

1. Use the lsdev command to ensure that the virtual SCSI adapter is available. For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   ent3     Available  Virtual I/O Ethernet Adapter (l-lan)
   vhost0   Available  Virtual SCSI Server Adapter
   vhost1   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   vtscsi0  Available  Virtual Target Device - Logical Volume
   vtscsi1  Available  Virtual Target Device - File-backed Disk
   vtscsi2  Available  Virtual Target Device - File-backed Disk
   ```
2. To create a virtual target device, which maps the virtual SCSI server adapter to a file or logical volume, run the mkbdsp command:

   ```
   mkbdsp -sp StoragePool -bd BackingDevice -vadapter VirtualSCSIServerAdapter -tn TargetDeviceName
   ```

   Where:
   - *StoragePool* is the name of the storage pool that contains the file or logical volume to which you plan to map the virtual SCSI server adapter. For example, fbPool.
   - *BackingDevice* is the name of the file or logical volume to which you plan to map the virtual SCSI server adapter. For example, devFile.
   - *VirtualSCSIServerAdapter* is the name of the virtual SCSI server adapter. For example, vhost4.
   - *TargetDeviceName* is the name of the target device. For example, fbvtd1.

   The storage is available to the client logical partition either the next time it starts, or the next time the appropriate virtual SCSI client adapter is probed (on a Linux logical partition), or configured (on an AIX logical partition).
3. View the newly created virtual target device by running the lsdev command. For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   vhost4   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   fbvtd1   Available  Virtual Target Device - File-backed Disk
   ```
4. View the logical connection between the newly created devices by running the lsmap command. For example, running `lsmap -vadapter vhost4` returns results similar to the following:

   ```
   SVSA     Physloc                    Client PartitionID
   ------------------------------------------------------
   vhost4   U9117.570.10C8BCE-V6-C2     0x00000000

   VTD             fbvtd1
   Status          Available
   LUN             0x8100000000000000
   Backing device  /var/vio/storagepools/fbPool/devFile
   Physloc
   ```

   The physical location is a combination of the slot number, in this case 2, and the logical partition ID. The virtual device can now be attached from the client logical partition.

If you later need to remove the virtual target device, you can do so by using the rmbdsp command.

**Creating a virtual target device on the Virtual I/O Server that maps to a file-backed virtual optical device:**

You can create a virtual target device on the Virtual I/O Server that maps the virtual SCSI adapter to a file-backed virtual optical device.

The following procedure can be repeated to provide additional virtual disk storage to any client logical partition.

Before you start, complete the following steps:

1. Ensure that the Virtual I/O Server is at version 1.5 or later.
2. Ensure that the virtual adapters for the Virtual I/O Server and the client logical partitions are created. This usually occurs during the creation of the logical partition profile. For information about creating the logical partition, see "Installing the Virtual I/O Server and client logical partitions" on page 61.

To create a virtual target device that maps a virtual SCSI server adapter to a file-backed virtual optical device, complete the following steps:

1. Use the lsdev command to ensure that the virtual SCSI adapter is available. For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   ent3     Available  Virtual I/O Ethernet Adapter (l-lan)
   vhost0   Available  Virtual SCSI Server Adapter
   vhost1   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   vtscsi0  Available  Virtual Target Device - Logical Volume
   vtscsi1  Available  Virtual Target Device - File-backed Disk
   vtscsi2  Available  Virtual Target Device - File-backed Disk
   ```

2. To create a virtual target device, which maps the virtual SCSI server adapter to a file-backed virtual optical device, run the mkvdev command:

   ```
   mkvdev -fbo -vadapter VirtualSCSIServerAdapter
   ```

   where *VirtualSCSIServerAdapter* is the name of the virtual SCSI server adapter. For example, vhost1.

   **Note:** No backing device is specified when creating virtual target devices for file-backed virtual optical devices because the drive is considered to contain no media. For information about loading media into a file-backed optical drive, see the loadopt command.

   The optical device is available to the client logical partition either the next time it starts, or the next time the appropriate virtual SCSI client adapter is probed (on a Linux logical partition), or configured (on an AIX logical partition).

3. View the newly created virtual target device by running the lsdev command. For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   vhost4   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   vtopt0   Available  Virtual Target Device - File-backed Optical
   ```

4. View the logical connection between the newly created devices by running the lsmap command. For example, running `lsmap -vadapter vhost1` returns results similar to the following:

   ```
   SVSA     Physloc                  Client PartitionID
   ----------------------------------------------------
   vhost1   U9117.570.10C8BCE-V6-C2  0x00000000

   VTD               vtopt0
   LUN               0x8200000000000000
   Backing device    Physloc
   ```

   The physical location is a combination of the slot number, in this case 2, and the logical partition ID. The virtual device can now be attached from the client logical partition.

You can use the loadopt command to load file-backed virtual optical media into the file-backed virtual optical device.

If you later need to remove the virtual target device, you can do so by using the rmvdev command.

## Creating logical volume storage pools on the Virtual I/O Server

You can create a logical volume storage pool on the Virtual I/O Server using the mksp and mkbdsp commands.

Before you start, ensure that the Virtual I/O Server is at version 1.5 or later.

Logical volume storage pools are volume groups, which are collections of one or more physical volumes. The physical volumes that comprise a logical volume storage pool can be of varying sizes and types.

To create a logical volume storage pool, complete the following steps:

1. Create a logical volume storage pool by running the mksp command:

   ```
   mksp -f dev_clients hdisk2 hdisk4
   ```

   In this example, the name of the storage pool is `dev_clients` and it contains `hdisk2` and `hdisk4`.
2. Define a logical volume, which will be visible as a disk to the client logical partition. The size of this logical volume will act as the size of disks that will be available to the client logical partition. Use the mkbdsp command to create a 11 GB logical volume called `dev_dbsrv` as follows:

   ```
   mkbdsp -sp dev_clients 11G -bd dev_dbsrv
   ```

   If you also want to create a virtual target device, which maps the virtual SCSI server adapter to the logical volume, add `-vadapter vhost`*`x`* to the end of the command. For example:

   ```
   mkbdsp -sp dev_clients 11G -bd dev_dbsrv -vadapter vhost4
   ```

## Creating file storage pools on the Virtual I/O Server

You can create a file storage pool on the Virtual I/O Server using the mksp and mkbdsp commands.

Before you start, ensure that the Virtual I/O Server is at version 1.5 or later.

File storage pools are created within a parent logical volume storage pool and contain a logical volume containing a filesystem with files.

To create a file storage pool, complete the following steps:

1. Create a file storage pool by running the mksp command:

   ```
   mksp -fb dev_fbclt -sp dev_clients -size 7g
   ```

   In this example, the name of the file storage pool is `dev_fbclt` and the parent storage pool is `dev_clients`.
2. Define a file, which will be visible as a disk to the client logical partition. The size of the file will act as the size of disks that will be available to the client logical partition. Use the mkbdsp command to create a 3 GB file called `dev_dbsrv` as follows:

   ```
   mkbdsp -sp dev_fbclt 3G -bd dev_dbsrv
   ```

   If you also want to create a virtual target device, which maps the virtual SCSI server adapter to the file, add `-vadapter vhost`*`x`* to the end of the command. For example:

   ```
   mkbdsp -sp dev_fbclt 3G -bd dev_dbsrv -vadapter vhost4
   ```

## Creating the virtual media repository on the Virtual I/O Server

You can create the virtual media repository on the Virtual I/O Server using the mkrep command.

Before you start, ensure that the Virtual I/O Server is at version 1.5 or later.

The virtual media repository provides a single container to store and manage file-backed virtual optical media files. Media stored in the repository can be loaded into file-backed virtual optical devices for exporting to client partitions.

Only one repository can be created within a Virtual I/O Server.

To create the virtual media repository, run the mkrep command:

```
mkrep -sp prod_store -size 6g
```

In this example, the name of the parent storage pool is `prod_store`.

## Creating volume groups and logical volumes on the Virtual I/O Server

You can create logical volumes and volume groups on the Virtual I/O Server using the mkvg and mklv commands.

To create a logical volume, use the mklv command. To create the logical volume on a separate disk, you must first create a volume group and assign one or more disks by using the mkvg command.

1. Create a volume group and assign a disk to this volume group by using the mkvg command. In this example, the name of the volume group is `rootvg_clients`

   ```
   mkvg -f -vg rootvg_clients hdisk2
   ```

2. Define a logical volume, which will be visible as a disk to the client logical partition. The size of this logical volume will act as the size of disks that will be available to the client logical partition. Use the mklv command to create a 2 GB logical volume as follows:

   ```
   mklv -lv rootvg_dbsrv rootvg_clients 2G
   ```

## Configure the Virtual I/O Server to support SCSI-2 reserve functions

Understand the virtual SCSI setup requirements to support applications using SCSI reserve and release.

Virtual I/O Server versions 1.3 and later provide support for applications that are enabled to use SCSI-2 reserve functions that are controlled by the client logical partition. Typically, SCSI reserve and release is used in clustered environments where contention for SCSI disk resources might require greater control. To ensure that Virtual I/O Server supports these environments, configure the Virtual I/O Server to support SCSI-2 reserve and release. If the applications you are using provide information about the policy to use for the SCSI-2 reserve functions on the client logical partition, follow those procedures for setting the reserve policy.

Complete the following tasks to configure the Virtual I/O Server to support SCSI-2 reserve environments:

1. Configure the Virtual I/O Server reserve_policy for single_path, using the following command:

   ```
   chdev -dev1 hdiskN -attr reserve_policy=single_path
   ```

   **Note:** Perform this task when the device is not in use. If you run this command while the device is open or in use, then you must use the **-perm** flag with this command. If you use the **-perm** flag, the changes do not take effect until the device is unconfigured and reconfigured.

2. Configure the client_reserve feature on the Virtual I/O Server.
   - If you are creating a virtual target device, use the following command:

     ```
     mkvdev -vdev hdiskN -vadapter vhostN -attr client_reserve=yes
     ```

     where *hdiskN* is the virtual target device name and *vhostN* is the virtual SCSI server adapter name.
   - If the virtual target device has already been created, use the following command:

     ```
     chdev -dev vtscsiN -attr client_reserve=yes
     ```

     where *vtscsiN* is the virtual device name.

3. On the Virtual client, complete the following steps to configure the SCSI reserve and release support for the virtual disk backed by the physical disk that you configured in step 1:
   a. Set the reserve policy on the Virtual client to single_path, using the following command:

      ```
      chdev -a reserve_policy=single_path -1 hdiskN
      ```

where *hdiskN* is the virtual disk name

**Note:** Perform this task when the device is not in use. If you run this command while the device is open or in use, then you must use the **-p** flag. In that case, the changes do not take effect until the device is unconfigured and reconfigured.

b. Set the hcheck_cmd attribute so that the MPIO code uses the inquiry option. If the hcheck_cmd attribute is set to **test unit ready** and the backing device is reserved, then *test unit ready* will fail and log an error on the client.

```
chdev -a hcheck_cmd=inquiry -1 hdiskN
```

where *hdiskN* is the virtual disk name.

## Identifying exportable disks

To export a physical volume as a virtual device, the physical volume must have an IEEE volume attribute, a unique identifier (UDID), or a physical identifier (PVID).

To identify exportable disks, complete the following steps:

1. Determine whether a device has an IEEE volume attribute identifier by running the following command from the Virtual I/O Server command line:

```
lsdev -dev hdiskX -attr
```

Disks with an IEEE volume attribute identifier have a value in the `ieee_volname` field. Output similar to the following is displayed:

```
...
cache_method   fast_write                       Write Caching method
   False
ieee_volname   600A0B800012DD0D00000AB441ED6AC IEEE Unique volume name
   False
lun_id         0x001a000000000000               Logical Unit Number
   False
...
```

If the `ieee_volname` field does not appear, then the device does not have an IEEE volume attribute identifier.

2. If the device does not have an IEEE volume attribute identifier, then determine whether the device has a UDID by completing the following steps:

a. Type `oem_setup_env`.

b. Type `odmget -qattribute=unique_id CuAt`. The disks that have a UDID are listed. Output similar to the following is displayed:

```
CuAt:
 name = "hdisk1"
 attribute = "unique_id"
 value = "2708ECVBZ1SC10IC35L146UCDY10-003IBMscsi"
 type = "R"
 generic = ""
 rep = "nl"
 nls_index = 79

CuAt:
 name = "hdisk2"
 attribute = "unique_id"
 value = "210800038FB50AST373453LC03IBMscsi"
 type = "R"
 generic = ""
 rep = "nl"
 nls_index = 79
```

Devices in the list that are accessible from other Virtual I/O Server partitions can be used in virtual SCSI MPIO configurations.

c. Type `exit`.

3. If the device does not have either an IEEE volume attribute identifier or a UDID, then determine whether the device has a PVID by running the following command:

```
lspv
```

The disks and their respective PVIDs are listed. Output similar to the following is displayed:

```
NAME        PVID                VG       STATUS
hdisk0      00c5e10c1608fd80    rootvg   active
hdisk1      00c5e10cf7eb2195    rootvg   active
hdisk2      00c5e10c44df5673    None
hdisk3      00c5e10cf3ba6a9a    None
hdisk4      none                None
```

4. If the device does not have either an IEEE volume attribute identifier, a UDID, or a PVID, then complete one of the following tasks to assign an identifier:

a. Upgrade your vendor software and then repeat this entire procedure, Identifying exportable disks, from the beginning. The latest versions of some vendor software include support for identifying devices using a UDID. Before upgrading, ensure that you preserve any virtual SCSI devices that you created when using the versions of the software that did not support identifying devices using a UDID. For information and upgrade instructions, see the documentation provided by your vendor software.

b. If the upgraded vendor software does not produce a UDID or IEEE volume attribute identifier, then put a PVID on the physical volume by running the following command:

```
chdev -dev hdiskX -attr pv=yes
```

## Configuring virtual Ethernet on the Virtual I/O Server

You can configure virtual Ethernet devices by deploying a system plan, create and configure a Shared Ethernet Adapter, and configure a Link Aggregation device.

### Creating a Shared Ethernet Adapter using HMC version 7

You can create a Shared Ethernet Adapter on the Virtual I/O Server so that client logical partitions can access the external network without needing to own a physical Ethernet adapter.

If you plan to use a Shared Ethernet Adapter with a Host Ethernet Adapter (or Integrated Virtual Ethernet), ensure that the Logical Host Ethernet Adapter (LHEA) on the Virtual I/O Server is set to promiscuous mode. For instructions, see “Setting the LHEA to promiscuous mode” on page 85.

To create a Shared Ethernet Adapter on the Virtual I/O Server using the Hardware Management Console (HMC), version 7 or later, complete the following steps:

1. In the navigation area, expand **Systems Management** → **Servers** and select the server on which the Virtual I/O Server logical partition is located.
2. In the contents are, select the Virtual I/O Server logical partition.
3. Click **Tasks** and select **Configuration** → **Manage Profiles**. The Managed Profiles page is displayed.
4. Select the profile in which you want to create the Shared Ethernet Adapter and click **Actions** → **Edit**. The Logical Partition Profile Properties page is displayed.
5. Click the **Virtual Adapters** tab.
6. Click **Actions** → **Create** → **Ethernet adapter**.
7. Select **IEEE 802.1Q-compatible adapter**.
8. If you are using multiple VLANs, add any additional VLAN IDs for the client logical partitions that must communicate with the external network using this virtual adapter.

9. Select **Access external network** to use this adapter as a gateway between VLANs and an external network. This Ethernet adapter is configured as part of the Shared Ethernet Adapter.
10. If you are not using Shared Ethernet Adapter failover, you can use the default trunk priority. If you are using Shared Ethernet Adapter failover, then set the trunk priority for the primary share Ethernet adapter to a lower number than that of the backup Shared Ethernet Adapter.
11. When you are finished, click **OK**.
12. Assign or create one of the following real adapters:
    - Assign a physical Ethernet adapter to the Virtual I/O Server.
    - If you plan to aggregate more than one physical Ethernet adapter into a Link Aggregation or EtherChannel device, then assign multiple physical Ethernet adapters to the Virtual I/O Server.
    - If you plan to use the Shared Ethernet Adapter with a Host Ethernet Adapter, then create an LHEA for the Virtual I/O Server logical partition.
13. Click **OK** to exit the Logical Partition Profile Properties page.
14. Click **Close** to exit the Managed Profiles page.
15. Repeat this procedure for additional Shared Ethernet Adapters that you require.

When you are finished, configure the Shared Ethernet Adapter using the Virtual I/O Server command-line interface. For instructions, see “Configuring a Shared Ethernet Adapter” on page 86.

**Setting the LHEA to promiscuous mode:**

To use a Shared Ethernet Adapter with a Host Ethernet Adapter (or Integrated Virtual Ethernet), you must set the Logical Host Ethernet Adapter (LHEA) to promiscuous mode.

Before you start, use the Hardware Management Console (HMC) to determine the physical port of the Host Ethernet Adapter that is associated with the Logical Host Ethernet port. Determine this information for the Logical Host Ethernet port that is the real adapter of the Shared Ethernet Adapter on the Virtual I/O Server. You can find this information in the partition properties of the Virtual I/O Server, and the managed system properties of the server on which the Virtual I/O Server is located.

To set the Logical Host Ethernet port (that is the real adapter of the Shared Ethernet Adapter) to promiscuous mode, complete the following steps using the HMC:

1. In the navigation area, expand **Systems Management** and click **Servers**.
2. In the contents area, select the server on which the Virtual I/O Server logical partition is located.
3. Click **Tasks** and select **Hardware (information)** → **Adapters** → **Host Ethernet**. The HEAs page is shown.
4. Select the physical location code of the Host Ethernet Adapter.
5. Select the physical port associated with the Logical Host Ethernet port on the Virtual I/O Server logical partition, and click **Configure**. The HEA Physical Port Configuration page is shown.
6. Select **VIOS** in the Promiscuous LPAR field.
7. Click **OK** twice to return to the contents area.

## Creating a Shared Ethernet Adapter using HMC version 6

You can create a Shared Ethernet Adapter on the Virtual I/O Server so that client logical partitions can access the external network without needing to own a physical Ethernet adapter.

To create a Shared Ethernet Adapter on the Virtual I/O Server using the Hardware Management Console (HMC), version 6 or earlier, complete the following steps:

1. On the HMC, right-click the profile for the Virtual I/O Server and select **Properties**.
2. Create a virtual Ethernet adapter using the **Virtual I/O** tab by choosing **Ethernet** in the Create Adapters area.

3. On the Virtual Ethernet Adapter Properties tab, choose the slot number for the virtual adapter and PVID (this PVID will be the default ID used later). Select **Trunk Adapter** to use this adapter as a gateway between VLANs and an external network. This Ethernet adapter is configured as part of the Shared Ethernet Adapter.
4. Select the **IEEE 802.1Q-compatible adapter** check box.
5. If you are using multiple VLANs, add any additional VLAN IDs for the client logical partitions that must communicate with the external network using this virtual adapter.
6. Repeat this procedure for additional Shared Ethernet Adapters that you require.

When you are finished, configure the Shared Ethernet Adapter using the Virtual I/O Server command-line interface. For instructions, see "Configuring a Shared Ethernet Adapter."

## Configuring a Shared Ethernet Adapter

Find instructions for configuring Shared Ethernet Adapters.

Before you can configure a Shared Ethernet Adapter, you must first create the adapter using the Hardware Management Console (HMC). For instructions, see one of the following tasks:

- "Creating a Shared Ethernet Adapter using HMC version 7" on page 84
- "Creating a Shared Ethernet Adapter using HMC version 6" on page 85

To configure a Shared Ethernet Adapter using the Virtual I/O Server, complete the following steps:

1. Verify that the virtual Ethernet trunk adapter is available by running the following command:
   ```
   lsdev -virtual
   ```
2. Identify the appropriate physical Ethernet adapter that will be used to create the Shared Ethernet Adapter by running the following command:
   ```
   lsdev -type adapter
   ```

   **Notes:**
   - Ensure that TCP/IP is not configured on the interface for the physical Ethernet adapter. If TCP/IP is configured, the mkvdev command in the next step fails.
   - You can also use a Link Aggregation, or EtherChannel, device as the Shared Ethernet Adapter.
   - If you plan to use the Host Ethernet Adapter or Integrated Virtual Ethernet with the Shared Ethernet Adapter, ensure that you use the Logical Host Ethernet Adapter to create the Shared Ethernet Adapter.
3. Configure the Shared Ethernet Adapter by running the following command:
   ```
   mkvdev -sea target_device -vadapter virtual_ethernet_adapters \
   -default DefaultVirtualEthernetAdapter -defaultid SEADefaultPVID
   ```

   Where:

   *target_device*
   The physical adapter being used as part of the Shared Ethernet Adapter device.

   *virtual_ethernet_adapters*
   The virtual Ethernet adapter or adapters that will use the Shared Ethernet Adapter.

   *DefaultVirtualEthernetAdapter*
   The default virtual Ethernet adapter used to handle untagged packets. If you have only one virtual Ethernet adapter for this logical partition, use it as the default.

   *SEADefaultPVID*
   The PVID associated with your default virtual Ethernet adapter.

For example, to create Shared Ethernet Adapter ent3 with ent0 as the physical Ethernet adapter (or Link Aggregation) and ent2 as the only virtual Ethernet adapter (defined with a PVID of 1), type the following command:

```
mkvdev -sea ent0 -vadapter ent2 -default ent2 -defaultid 1
```

4. Verify that the Shared Ethernet Adapter was created by running the following command:

   ```
   lsdev -virtual
   ```

5. Do you plan to access the Virtual I/O Server from the network with the physical device used to create the Shared Ethernet Adapter?
   - Yes: Go to step 6.
   - No: You are finished with this procedure and do not need to complete the remaining steps.
6. Do you plan to define IP addresses on any VLANs other than the VLAN specified by the PVID of the Shared Ethernet Adapter?
   - Yes: Go to step 7 to create VLAN pseudo-devices.
   - No: Go to step 8 to configure a TCP/IP connection.
7. To configure VLAN pseudo-devices, complete the following steps:
   a. Create a VLAN pseudo-device on the Shared Ethernet Adapter by running the following command:

      ```
      mkvdev -vlan TargetAdapter -tagid TagID
      ```

      Where:
      - *TargetAdapter* is the Shared Ethernet Adapter.
      - *TagID* is the VLAN ID that you defined when creating the virtual Ethernet adapter associated with the Shared Ethernet Adapter.

      For example, to create a VLAN pseudo-device using the Shared Ethernet Adapter ent3 that you just created with a VLAN ID of 1, type the following command:

      ```
      mkvdev -vlan ent3 -tagid 1
      ```

   b. Verify that the VLAN pseudo-device was created by running the following command:

      ```
      lsdev -virtual
      ```

   c. Repeat this step for any additional VLAN pseudo-devices that you need.
8. Run the following command to configure the first TCP/IP connection. The first connection must be on the same VLAN and logical subnet as the default gateway.

   ```
   mktcpip -hostname Hostname -inetaddr Address -interface Interface -netmask \
   SubnetMask -gateway Gateway -nsrvaddr NameServerAddress -nsrvdomain Domain
   ```

   Where:
   - *Hostname* is the host name of the Virtual I/O Server
   - *Address* is the IP address you want to use for the TCP/IP connection
   - *Interface* is the interface associated with either the Shared Ethernet Adapter device or a VLAN pseudo-device. For example, if the Shared Ethernet Adapter device is ent3, the associated interface is en3.
   - *Subnetmask* is the subnet mask address for your subnet.
   - *Gateway* is the gateway address for your subnet.
   - *NameServerAddress* is the address of your domain name server.
   - *Domain* is the name of your domain.

   If you do not have additional VLANs, then you are finished with this procedure and do not need to complete the remaining step.
9. Run the following command to configure additional TCP/IP connections:

   ```
   chdev -dev interface -perm -attr netaddr=IPaddress -attr netmask=netmask
   -attr state=up
   ```

When using this command, enter the interface (en*X*) associated with either the Shared Ethernet Adapter device or VLAN pseudo-device.

The Shared Ethernet Adapter is now configured. After you configure the TCP/IP connections for the virtual adapters on the client logical partitions using the client logical partitions' operating systems, those logical partitions can communicate with the external network.

**Related concepts**

"Shared Ethernet Adapter failover" on page 58
Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server logical partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption.

## Configuring a Link Aggregation or EtherChannel device

Configure a Link Aggregation device, also called an EtherChannel device, by using the mkvdev command. A Link Aggregation device can be used as the physical Ethernet adapter in the Shared Ethernet Adapter configuration.

Configure a Link Aggregation device by typing the following command:

```
mkvdev -lnagg TargetAdapter ... [-attr Attribute=Value ...]
```

For example, to create Link Aggregation device `ent5` with physical Ethernet adapters `ent3`, `ent4`, and backup adapter `ent2`, type the following:

```
mkvdev -lnagg ent3,ent4 -attr backup_adapter=ent2
```

After the Link Aggregation device is configured, you can add adapters to it, remove adapters from it, or modify its attributes using the cfglnagg command.

# Configuring the Tivoli agents and clients on the Virtual I/O Server

You can configure and start the Tivoli Monitoring agent, Tivoli Usage and Accounting Manager, the Tivoli Storage Manager client, and the Tivoli TotalStorage Productivity Center agents.

## Configuring the Tivoli Monitoring agent

You can configure and start the Tivoli Monitoring agent on the Virtual I/O Server.

With Tivoli Monitoring System Edition , you can monitor the health and availability of multiple servers (including the Virtual I/O Server) from the Tivoli Enterprise Portal. Tivoli Monitoring System Edition gathers data from the Virtual I/O Server, including data about physical volumes, logical volumes, storage pools, storage mappings, network mappings, real memory, processor resources, mounted file system sizes, and so on. From the Tivoli Enterprise Portal, you can view a graphical representation of the data, use predefined thresholds to alert you on key metrics, and resolve issues based on recommendations provided by the Expert Advice feature of Tivoli Monitoring.

Before you start, complete the following tasks:

- Ensure that the Virtual I/O Server is running fix pack 8.1.0.
- Verify that you are a super administrator of the HMC.
- Verify that you are the prime administrator of the Virtual I/O Server.

To configure and start the monitoring agent, complete the following steps:

1. List all of the available monitoring agents using the **lssvc** command. For example,
   ```
   $lssvc
   ITM_base
   ```
2. Based on the output of the **lssvc** command, decide which monitoring agent you want to configure. For example, `ITM_base`

3. List all of the attributes that are associated with the monitoring agent using the **cfgsvc** command. For example:

   ```
   $cfgsvc –ls ITM_base
    HOSTNAME
   RESTART_ON_REBOOT
   MANAGING_SYSTEM
   ```

4. Configure the monitoring agent with its associated attributes using the **cfgsvc** command:

   ```
   cfgsvc ITM_agent_name -attr Restart_On_Reboot=value hostname=name_or_address1
   managing_system=name_or_address2
   ```

   Where:

   - *ITM_agent_name* is the name of the monitoring agent. For example, ITM_base.
   - *value* must be either TRUE of FALSE as follows:
     - TRUE: *ITM_agent_name* restarts whenever the Virtual I/O Server restarts
     - FALSE: *ITM_agent_name* does not restart whenever the Virtual I/O Server restarts
   - *name_or_address1* is either the hostname or IP address of the Tivoli Enterprise Monitoring Server (TEMS) server to which *ITM_agent_name* sends data.
   - *name_or_address2* is either the hostname of IP address of the Hardware Management Console (HMC) attached to the managed system on which the Virtual I/O Server with the monitoring agent is located.

   For example:

   ```
   cfgsvc ITM_base –attr Restart_On_Reboot=TRUE hostname=tems_server managing_system=hmc_console
   ```

   In this example, the ITM_base monitoring agent is configured to send data to tems_server, and to restart whenever the Virtual I/O Server restarts.

5. Start the monitoring agent using the **startsvc** command. For example:

   ```
   startsvc ITM_base
   ```

6. From the HMC, complete the following steps so that the monitoring agent can gather information from the HMC.

   **Note:** After you configure a secure shell connection for one monitoring agent, you do not need to configure it again for any additional agents.

   a. Determine the name of the managed system on which the Virtual I/O Server with the monitoring agent is located.

   b. Obtain the public key for the Virtual I/O Server by running the following command:

      ```
      viosvrcmd -m managed_system_name -p vios_name -c "cfgsvc -key ITM_agent_name"
      ```

      Where:

      - *managed_system_name* is the name of the managed system on which the Virtual I/O Server with the monitoring agent or client is located.
      - *vios_name* is the name of the Virtual I/O Server logical partition (with the monitoring agent) as defined on the HMC.
      - *ITM_agent_name* is the name of the monitoring agent. For example, ITM_base.

   c. Update the authorized_key2 file on the HMC by running the mkauthkeys command:

      ```
      mkauthkeys --add public_key
      ```

      where *public_key* is the output from the viosvrcmd command in step 6b.

   For example:

   ```
   $ viosvrcmd -m commo126041 -p VIOS7 -c "cfgsvc ITM_base -key"
   ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAvjDZ
     sS0guWzfzfp9BbweG0QMXv1tbDrtyWsgPbA2ExHA+xduWA51K0oFGarK2F
   ```

```
  C7e7NjKW+UmgQbrh/KSyKKwozjp4xWGNGhLmfan85ZpFR7wy9UQG1bLgXZ
  xYrY7yyQQQODjvwosWAfzkjpG3iW/xmWD5PKLBmob2QkKJbxjne+wqGwHT
  RYDGIiyhCBIdfFaLZgkXTZ2diZ98rL8LIv3qb+TsM1B28AL4t+1OGGeW24
  2lsB+8p4kamPJCYfKePHo67yP4NyKyPBFHY3TpTrca4/y1KEBT0Va3Pebr
  5JEIUvWYs6/RW+bUQk1Sb6eYbcRJFHhN5l3F+ofd0vj39zwQ== root@vi
  os7.vios.austin.ibm.com
$ mkauthkeys --add 'ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAvjDZ
  sS0guWzfzfp9BbweG0QMXv1tbDrtyWsgPbA2ExHA+xduWA51K0oFGarK2F
  C7e7NjKW+UmgQbrh/KSyKKwozjp4xWGNGhLmfan85ZpFR7wy9UQG1bLgXZ
  xYrY7yyQQQODjvwosWAfzkjpG3iW/xmWD5PKLBmob2QkKJbxjne+wqGwHT
  RYDGIiyhCBIdfFaLZgkXTZ2diZ98rL8LIv3qb+TsM1B28AL4t+1OGGeW24
  2lsB+8p4kamPJCYfKePHo67yP4NyKyPBFHY3TpTrca4/y1KEBT0Va3Pebr
  5JEIUvWYs6/RW+bUQk1Sb6eYbcRJFHhN5l3F+ofd0vj39zwQ== root@vi
  os7.vios.austin.ibm.com'
```

When you are finished, you can view the data gathered by the monitoring agent from the Tivoli Enterprise Portal.

**Related information**

Tivoli Monitoring 6.1 documentation

Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide

## Configuring the Tivoli Usage and Accounting Manager agent

You can configure and start the Tivoli Usage and Accounting Manager agent on the Virtual I/O Server.

With Virtual I/O Server 1.4, you can install and configure the Tivoli Usage and Accounting Manager agent on the Virtual I/O Server. Tivoli Usage and Accounting Manager helps you track, allocate, and invoice your IT costs by collecting, analyzing, and reporting on the actual resources used by entities such as cost centers, departments, and users. Tivoli Usage and Accounting Manager can gather data from multi-tiered datacenters that include Windows, AIX, Virtual I/O Server, HP/UX Sun Solaris, Linux and VMware.

Before you start, ensure that the Virtual I/O Server is installed. The Tivoli Usage and Accounting Manager agent is packaged with the Virtual I/O Server and is installed when the Virtual I/O Server is installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 61.

To configure and start the Tivoli Usage and Accounting Manager agent, complete the following steps:

1. Optional: Add optional variables to the A_config.par file to enhance data collection. The A_config.par file is located at /home/padmin/tivoli/ituam/A_config.par.
2. List all of the available Tivoli Usage and Accounting Manager agents using the **lssvc** command. For example,

   ```
   $lssvc
   ITUAM_base
   ```
3. Based on the output of the **lssvc** command, decide which Tivoli Usage and Accounting Manager agent you want to configure. For example, `ITUAM_base`
4. List all of the attributes that are associated with the Tivoli Usage and Accounting Manager agent using the **cfgsvc** command. For example:

   ```
   $cfgsvc –ls ITUAM_base
    ACCT_DATA0
   ACCT_DATA1
   ISYSTEM
   IPROCESS
   ```
5. Configure the Tivoli Usage and Accounting Manager agent with its associated attributes using the **cfgsvc** command:

   ```
   cfgsvc ITUAM_agent_name -attr ACCT_DATA0=value1 ACCT_DATA1=value2 ISYSTEM=value3 IPROCESS=value4
   ```

   Where:

- *ITUAM_agent_name* is the name of the Tivoli Usage and Accounting Manager agent. For example, ITUAM_base.
- *value1* is the size (in MB) of the first data file that holds daily accounting information.
- *value2* is the size (in MB) of the second data file that holds daily accounting information.
- *value3* is the time (in minutes) when the agent generates system interval records.
- *value4* is the time (in minutes) when the system generates aggregate process records.

6. Start the Tivoli Usage and Accounting Manager agent using the **startsvc** command. For example:

   ```
   startsvc ITUAM_base
   ```

After you start the Tivoli Usage and Accounting Manager agent, it begins to collect data and generate log files. You can configure the Tivoli Usage and Accounting Manager server to retrieve the log files, which are then processed by the Tivoli Usage and Accounting Manager Processing Engine. You can work with the data from the Tivoli Usage and Accounting Manager Processing Engine as follows:

- You can generate customized reports, spreadsheets, and graphs. Tivoli Usage and Accounting Manager provides full data access and reporting capabilities by integrating Microsoft® SQL Server Reporting Services or Crystal Reports with a Database Management System (DBMS).
- You can view high-level and detailed cost and usage information.
- You can allocate, distribute, or charge IT costs to users, cost centers, and organizations in a manner that is fair, understandable, and reproducible.

For more information, see the Tivoli Usage and Accounting Manager Information Center.

## Configuring the Tivoli Storage Manager client

You can configure theTivoli Storage Manager client on the Virtual I/O Server.

With Virtual I/O Server 1.4, you can install and configure the Tivoli Storage Manager client on the Virtual I/O Server. With Tivoli Storage Manager, you can protect your data from failures and other errors by storing backup and disaster-recovery data in a hierarchy of offline storage. Tivoli Storage Manager can help protect computers running a variety of different operating environments, including the Virtual I/O Server, on a variety of different hardware. If you configure the Tivoli Storage Manager client on the Virtual I/O Server, you can include the Virtual I/O Server in your standard backup framework.

Before you start, ensure that the Virtual I/O Server is installed. The Tivoli Storage Manager client is packaged with the Virtual I/O Server and is installed when the Virtual I/O Server is installed. For instructions, see “Installing the Virtual I/O Server and client logical partitions” on page 61.

To configure and start the Tivoli Storage Manager client, complete the following steps:

1. List all of the available Tivoli Storage Manager clients using the **lssvc** command. For example,

   ```
   $lssvc
   TSM_base
   ```

2. Based on the output of the **lssvc** command, decide which Tivoli Storage Manager client you want to configure. For example, `TSM_base`
3. List all of the attributes that are associated with the Tivoli Storage Manager client using the **cfgsvc** command. For example:

   ```
   $cfgsvc –ls TSM_base
    SERVERNAME
   SERVERIP
   NODENAME
   ```

4. Configure the Tivoli Storage Manager client with its associated attributes using the **cfgsvc** command:

   ```
   cfgsvc TSM_client_name -attr SERVERNAME=hostname SERVERIP=name_or_address NODENAME=vios
   ```

   Where:

   - *TSM_client_name* is the name of the Tivoli Storage Manager client. For example, TSM_base.

- *hostname* is the host name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated.
- *name_or_address* is the IP address or domain name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated.
- *vios* is the name of the machine on which the Tivoli Storage Manager client is installed. The name must match the name registered on the Tivoli Storage Manager server.

5. Ask the Tivoli Storage Manager administrator to register the client node, the Virtual I/O Server, with the Tivoli Storage Manager server.

After you are finished, you are ready to back up and restore the Virtual I/O Server using the Tivoli Storage Manager. For instructions, see the following procedures:

- "Backing up the Virtual I/O Server using Tivoli Storage Manager" on page 108
- "Restoring the Virtual I/O Server using Tivoli Storage Manager" on page 113

## Installing and configuring the TotalStorage Productivity Center agents

You can install, configure, and start the TotalStorage Productivity Center agents on the Virtual I/O Server.

With Virtual I/O Server 1.5.2, you can install and configure the TotalStorage Productivity Center agents on the Virtual I/O Server. TotalStorage Productivity Center is an integrated, storage infrastructure management suite that is designed to help simplify and automate the management of storage devices, storage networks, and capacity utilization of file systems and databases. When you install and configure the TotalStorage Productivity Center agents on the Virtual I/O Server, you can use the TotalStorage Productivity Center user interface to collect and view information about the Virtual I/O Server.

Before you start, complete the following tasks:

1. Use the ioslevel command to verify that the Virtual I/O Server is at version 1.5.2, or later.
2. Ensure that there are no other operations running on the Virtual I/O Server. Installing the TotalStorage Productivity Center consumes all of the processing time.
3. In addition to the memory required by the Virtual I/O Server logical partition, ensure that you have allocated a minimum of 1 GB of memory to the Virtual I/O Server for the TotalStorage Productivity Center agents.

To configure and start the TotalStorage Productivity Center agents, complete the following steps:

1. List all of the available TotalStorage Productivity Center agents using the **lssvc** command. For example,

   ```
   $lssvc
   TPC
   ```

   The TPC agent includes both the TPC_data and TPC_fabric agents. When you install and configure the TPC agent, you install and configure both the TPC_data and TPC_fabric agents.
2. List all of the attributes that are associated with the TotalStorage Productivity Center agent using the **lssvc** command. For example:

   ```
   $lssvc TPC
   A:
   S:
   devAuth:
   caPass:
   caPort:
   amRegPort:
   amPubPort:
   dataPort:
   devPort:
   newCA:
   oldCA:
   daScan:
   ```

```
daScript:
daInstall:
faInstall:
U:
```

The A, S, devAuth, and caPass attributes are required. The remainder of the attributes are optional. For more information about the attributes, see "Configuration attributes for Tivoli agents and clients" on page 130.

3. Install and configure the TotalStorage Productivity Center agent with its associated attributes using the **cfgsvc** command:

   ```
   cfgsvc TPC -attr S=tpc_server_hostname A=agent_manager_hostname devAuth=password_1 caPass=password_2
   ```

   Where:
   - *tpc_server_hostname* is the host name or IP address of the TotalStorage Productivity Center server that is associated with the TotalStorage Productivity Center agent.
   - *agent_manager_hostname* is the name or IP address of the Agent Manager.
   - *password_1* is the password required to authenticate to the TotalStorage Productivity Center device server.
   - *password_2* is the password required to authenticate to the common agent.
4. Select the language that you want to use during the installation and configuration.
5. Accept the license agreement to install the agents according to the attributes specified in step 3.
6. Start each TotalStorage Productivity Center agent using the **startsvc** command:
   - To start the TPC_data agent, run the following command:

     ```
     startsvc TPC_data
     ```
   - To start the TPC_fabric agent, run the following command:

     ```
     startsvc TPC_fabric
     ```

After you start the TotalStorage Productivity Center agents, you can perform the following tasks using the TotalStorage Productivity Center user interface:

1. Run a discovery job for the agents on the Virtual I/O Server.
2. Run probes, scans, and ping jobs to collect storage information about the Virtual I/O Server.
3. Generate reports using the Fabric Manager and the Data Manager to view the storage information gathered.
4. View the storage information gathered using the topology Viewer.

For more information, see the *TotalStorage Productivity Center support for agents on a Virtual I/O Server* PDF. To view or download the PDF, go to the TotalStorage Productivity Center v3.3.1.81 Interim Fix Web site.

## Configuring the Virtual I/O Server as an LDAP client

Virtual I/O Server version 1.4 can be configured as an LDAP client and then you can manage Virtual I/O Server from an LDAP server.

Before you start, gather the following information:

- The name of the Lightweight Directory Access Protocol (LDAP) server or servers to which you want the Virtual I/O Server to be an LDAP client.
- The administrator distinguish name (DN) and password for the LDAP server or servers to which you want the Virtual I/O Server to be an LDAP client.

To configure the Virtual I/O Server as an LDAP client, complete the following steps:

1. Change Virtual I/O Server users to LDAP users by running the following command:

   ```
   chuser -ldap username
   ```

where *username* is the name of the user you want to change to an LDAP user.

2. Set up the LDAP client by running the following command:

   ```
   mkldap –host ldapserv1 –bind cn=admin –passwd adminpwd
   ```

   Where:

   - *ldapserv1* is the LDAP server or list of LDAP servers to which you want the Virtual I/O Server to be an LDAP client
   - *cn=admin* is the administrator DN of *ldapserv1*
   - *adminpwd* is the password for *cn=admin*

   Configuring the LDAP client automatically starts communication between the LDAP server and the LDAP client (the Virtual I/O Server). To stop communication, use the stopnetsvc command.

# Managing the Virtual I/O Server

You can manage virtual SCSI and virtual Ethernet devices on the Virtual I/O Server, as well as back up, restore, update, and monitor the Virtual I/O Server.

Most of the information in this topic is specific to management in an HMC environment. For information about management tasks in an Integrated Virtualization Manager environment, see Chapter 4, “Integrated Virtualization Manager,” on page 163.

## Managing storage

You can import and export volume groups and storage pools, map virtual disks to physical disks, increase virtual SCSI device capacity, change the virtual SCSI queue depth, back up and restore files and file systems, and collect and view information using the TotalStorage Productivity Center.

### Importing and exporting volume groups and logical volume storage pools

You can use the importvg and exportvg commands to move a user-defined volume group from one system to another.

Consider the following when importing and exporting volume groups and logical volume storage pools:

- The import procedure introduces the volume group to its new system.
- You can use the importvg command to reintroduce a volume group or logical volume storage pool to the system that it had been previously associated with and had been exported from.
- The importvg command changes the name of an imported logical volume if a logical volume of that name already exists on the new system. If the importvg command must rename a logical volume, it prints an error message to standard error.
- The export procedure removes the definition of a volume group from a system.
- You can use the importvg and exportvg commands to add a physical volume that contains data to a volume group by putting the disk to be added in its own volume group.
- The rootvg volume group cannot be exported or imported.

**Importing volume groups and logical volume storage pools:**

You can use the importvg command to import a volume group or logical volume storage pool.

To import a volume group or logical volume storage pool, complete the following steps:

1. Run the following command to import the volume group or logical volume storage pool:

   ```
   importvg -vg volumeGroupName physicalVolumeName
   ```

   Where:

- *volumeGroupName* is an optional parameter that specifies the name to use for the imported volume group.
- *physicalVolumeName* is the name of a physical volume that belongs to the imported volume group.

2. If you know that the imported volume group or logical volume storage pool is not the parent of the virtual media repository or any file storage pools, then you are finished importing the volume group or logical volume storage pool and do not need to complete the remaining steps.
3. If you know that imported volume group or logical volume storage pool is the parent of the virtual media repository or any file storage pools, or if you are unsure, then complete the following steps:
   a. Run the `mount all` command to mount any file systems contained in the imported volume group or logical volume storage pool. This command might return errors for file systems that are already mounted.
   b. If you are importing a volume group or logical volume storage to the same system from which you exported it, run the `cfgdev` to reconfigure any devices that were unconfigured when you exported the volume group or logical volume storage pool.

To export a volume group or logical volume storage pool, see “Exporting volume groups and logical volume storage pools.”

**Exporting volume groups and logical volume storage pools:**

You can use the exportvg command to export a volume group or logical volume storage pool.

Before you start, complete the following tasks:

1. Determine whether the volume group or logical volume storage pool that you plan to export is a parent to the virtual media repository or to any file storage pools by completing the following steps:
   a. Run the lsrep command to determine whether the volume group or logical volume storage pool that you plan to export is a parent of the virtual media repository. The Parent Pool field displays the parent volume group or logical volume pool of the virtual media repository.
   b. Run the following command to determine whether a file storage pool is a child of the volume group or logical volume pool that you plan to export:

      ```
      lssp -detail -sp FilePoolName
      ```

      The results list the parent volume group or logical volume storage pool of the file storage pool.
2. If the volume group or logical volume storage pool that you plan to export is a parent of the virtual media repository or a file storage pool, then complete the following steps.

*Table 25. Prerequisites steps if the volume group or logical volume storage pool is a parent of the virtual media repository or a file storage pool*

| Parent of Virtual Media Repository | Parent of a file storage pool |
|---|---|
| 1. Unload the backing device of each file-backed optical virtual target device (VTD) that has a media file loaded, by completing the following steps:<br>a. Retrieve a list of the file-backed optical VTDs by running the following command:<br>`lsmap -all -type file_opt`<br>b. For each device that shows a backing device, run the following command to unload the backing device:<br>`unloadopt -vtd` *VirtualTargetDevice*<br>2. Unmount the Virtual Media Repository file system by running the following command:<br>`unmount /var/vio/VMLibrary` | 1. Unconfigure the virtual target devices (VTDs) associated with the files contained in the file storage pools by completing the following steps:<br>a. Retrieve a list of VTDs by running the following command:<br>`lssp -bd -sp` *FilePoolName*<br>where *FilePoolName* is the name of a file storage pool that is a child of the volume group or logical volume storage pool that you plan to export.<br>b. For each file that lists a VTD, run the following command:<br>`rmdev -dev` *VirtualTargetDevice* `-ucfg`<br>2. Unmount the file storage pool by running the following command:<br>`unmount /var/vio/storagepools/`*FilePoolName*<br>where *FilePoolName* is the name of the file storage pool to be unmounted. |

To export the volume group or logical volume storage pool, run the following commands:

1. `deactivatevg` *VolumeGroupName*
2. `exportvg` *VolumeGroupName*

To import a volume group or logical volume storage pool, see "Importing volume groups and logical volume storage pools" on page 94.

## Mapping virtual disks to physical disks

Find instructions for mapping a virtual disk on a client logical partition to its physical disk on the Virtual I/O Server.

This procedure shows how to map a virtual SCSI disk on an AIX client logical partition to the physical device (disk or logical volume) on the Virtual I/O Server.

To map a virtual disk to a physical disk, you need the following information. This information is gathered during this procedure:

- Virtual device name
- Slot number of the virtual SCSI client adapter
- Logical unit number (LUN) of the virtual SCSI device
- Client logical partition ID

Follow these steps to map a virtual disk on an AIX client logical partition to its physical disk on the Virtual I/O Server:

1. Display virtual SCSI device information on the AIX client logical partition by typing the following command:

   ```
   lscfg -l devicename
   ```

   This command returns results similar to the following:

   ```
   U9117.570.1012A9F-V3-C2-T1-L810000000000  Virtual SCSI Disk Drive
   ```

2. Record the slot number, which is located in the output, following the card location label *C*. This identifies the slot number of the virtual SCSI client adapter. In this example, the slot number is 2.
3. Record the LUN, which is located in the output, following the LUN label *L*. In this example, the LUN is 810000000000.
4. Record the logical partition ID of the AIX client logical partition:
   a. Connect to the AIX client logical partition. For example, using Telnet.
   b. On the AIX logical partition, run the `uname -L` command.

      Your results should look similar to the following:

      ```
      2  fumi02
      ```

      The logical partition ID is the first number listed. In this example, the logical partition ID is 2. This number is used in the next step.
   c. Type `exit`.
5. If you have multiple Virtual I/O Server logical partitions running on your system, determine which Virtual I/O Server logical partition is serving the virtual SCSI device. Use the slot number of the client adapter that is linked to a Virtual I/O Server, and a server adapter. Use the HMC command line to list information about virtual SCSI client adapters in the client logical partition.

   Log in to the HMC, and from the HMC command line, type `lshwres` . Specify the managed console name for the **-m** parameter and the client logical partition ID for the **lpar_ids** parameter.

   **Note:**

   - The managed console name, which is used for the **-m** parameter, is determined by typing `lssyscfg -r sys -F name` from the HMC command line.
   - Use the client logical partition ID recorded in Step 4 for the **-lpar_ids** parameter.

   For example:

   ```
   lshwres -r virtualio --rsubtype scsi -m fumi --filter lpar_ids=2
   ```

   This example returns results similar to the following:

   ```
   lpar_name=fumi02,lpar_id=2,slot_num=2,state=null,adapter_type=client,remote_lpar_id=1,
   remote_lpar_name=fumi01,remote_slot_num=2,is_required=1,backing_devices=none
   ```

   Record the name of the Virtual I/O Server located in the **remote_lpar_name** field and slot number of the virtual SCSI server adapter, which is located in the **remote_lpar_id** field. In this example, the name of the Virtual I/O Server is fumi01 and the slot number of the virtual SCSI server adapter is 1.
6. Log in to the Virtual I/O Server.
7. List virtual adapters and devices on the Virtual I/O Server by typing the following command:

   ```
   lsmap -all
   ```
8. Find the virtual SCSI server adapter (vhost*X*) that has a slot ID that matches the remote slot ID recorded in Step 7. On that adapter, run the following command:

   ```
   lsmap -vadapter devicename
   ```
9. From the list of devices, match the LUN recorded in Step 4 with LUNs listed. This is the physical device.

## Increasing virtual SCSI device capacity

Increase the size of virtual SCSI disks.

As storage demands increase for virtual client logical partitions, you can add physical storage to increase the size of your virtual devices and allocate that storage to your virtual environment. You can increase the capacity of your virtual SCSI devices by increasing the size of physical or logical volumes. With

Virtual I/O Server version 1.3 and later, you can do this without disrupting client operations. To increase the size of files and logical volumes based on storage pools, the Virtual I/O Server must be at version 1.5 or later.

To increase virtual SCSI device capacity, complete the following steps:

1. Increase the size of the physical volumes, logical volumes, or files:
   - Physical volumes: Consult your storage documentation to determine whether your storage subsystem supports expanding the size of a logical unit number (LUN).
   - Logical volumes based on volume groups:
     a. Run the extendlv command. For example: `extendlv lv3 100M`. This example increases logical volume *lv3* by 100 MB.
     b. If there is no additional space in the logical volume, complete the following tasks:
        1) Increase the size of the volume group by completing one of the following steps:
           – Increase the size of the physical volumes. Consult your storage documentation for instructions.
           – Add physical volumes to a volume group by running the extendvg command. For example: `extendvg vg1 hdisk2`. This example adds physical volume *hdisk2* to volume group *vg1*.
        2) Allocate the increased volume to partitions by resizing logical volumes. Run the extendlv command to increase the size of a logical volume.
   - Logical volumes based on storage pools:
     a. Run the chbdsp command. For example:`chbdsp -sp lvPool -bd lv3 -size 100M`. This example increases logical volume *lv3* by 100 MB.
     b. If there is no additional space in the logical volume, complete the following tasks:
        1) Increase the size of the logical volume storage pool by completing one of the following steps:
           – Increase the size of the physical volumes. Consult your storage documentation for instructions.
           – Add physical volumes to the storage pool by running the chsp command. For example: `chsp -add -sp sp1 hdisk2`. This example adds physical volume *hdisk2* to storage pool *sp1*.
        2) Allocate the increased volume to partitions by resizing logical volumes. Run the chbdsp command to increase the size of a logical volume.
   - Files:
     a. Run the chbdsp command. For example:`chbdsp -sp fbPool -bd fb3 -size 100M`. This example increases file *fb3* by 100 MB.
     b. If there is no additional space in the file, increase the size of the file storage pool by running the chsp command. For example:`chsp -add -sp fbPool -size 100M`. This example increases file storage pool *fbPool* by 100MB.
     c. If there is no additional space in the file storage pool, increase the size of the parent storage pool by completing one of the following tasks:
        – Increase the size of the physical volumes. Consult your storage documentation for instructions.
        – Add physical volumes to the parent storage pool by running the chsp command. For example:`chsp -add -sp sp1 hdisk2`. This example adds physical volume *hdisk2* to storage pool *sp1*.
        – Increase the size of the file storage pool by running the chsp command.
2. If you are running Virtual I/O Server versions prior to 1.3, then you need to either reconfigure the virtual device (using the cfgdev command) or restart the Virtual I/O Server.
3. If you are running Virtual I/O Server version 1.3 or later, then restarting or reconfiguring a logical partition is not required to begin using the additional resources. If the physical storage resources have

been set up and properly allocated to the system as a system resource, as soon as the Virtual I/O Server recognizes the changes in storage volume, the increased storage capacity is available to the client logical partitions.

4. On the client logical partition, ensure that the operating system recognizes and adjusts to the new size. For example, if AIX is the operating system on the client logical partition, run the following command:

   ```
   chvg -g vg1
   ```

   In this example, AIX examines all the disks in volume group *vg1* to see if they have grown in size. For the disks that have grown in size, AIX attempts to add additional physical partitions to physical volumes. If necessary, AIX will determine proper 1016 multiplier and conversion to the big volume group.

   **Related information**

   chvg Command

## Changing the virtual SCSI queue depth

Increasing the virtual SCSI queue depth might provide performance improvements for some virtual configurations. Understand the factors involved in determining a change to the virtual SCSI queue depth value.

The virtual SCSI queue depth value determines how many requests the disk head driver will queue to the virtual SCSI client driver at any one time. For AIX and Linux client logical partitions, you can change this value from the default value of 3 to any value from 1 to 256. You modify this value using the chdev command.

Increasing this value might improve the throughput of the disk in specific configurations. However, several factors must be taken into consideration. These factors include the value of the queue-depth attribute for all of the physical storage devices on the Virtual I/O Server being used as a virtual target device by the disk instance on the client logical partition, and the maximum transfer size for the virtual SCSI client adapter instance that is the parent device for the disk instance.

For AIX and Linux client logical partitions, the maximum transfer size for virtual SCSI client adapters is set by the Virtual I/O Server, which determines the value based on the resources available on the server and the maximum transfer size set for the physical storage devices on that server. Other factors include the queue depth and maximum transfer size of other devices involved in mirrored-volume-group or Multipath I/O (MPIO) configurations. Increasing the queue depth for some devices might reduce the resources available for other devices on that same shared adapter and decrease the throughput for those devices.

To change the queue depth for an AIX or Linux client logical partition, on the client logical partition use the chdev command with the **queue_depth=value** attribute as in the following example:

```
chdev -1 hdiskN -a "queue_depth=value"
```

*hdiskN* represents the name of a physical volume and *value* is the value you assign between 1 and 256.

To view the current setting for the queue_depth value, from the client logical partition issue the following command:

```
lsattr -E1 hdiskN
```

## Backing up and restoring files and file systems

You can use the backup and restore commands to back up and restore individual files or entire file systems.

Backing up and restoring files and files systems can be useful for tasks, such as saving a file-backed device.

The following commands are used to back up and restore files and files systems.

*Table 26. Backup and restore commands and their descriptions*

| Command | Description |
|---|---|
| backup | Backs up files and file systems to media, such as physical tape and disk. For example:<br>• You can back up all the files and subdirectories in a directory using full path names or relative path names.<br>• You can back up the root file system.<br>• You can back up all the files in the root file system that have been modified since the last backup.<br>• You can back up virtual optical media files from the virtual media repository. |
| restore | Reads archives created by the backup command and extracts the files stored there. For example:<br>• You can restore a specific file into the current directory.<br>• You can restore a specific file from tape into the virtual media repository.<br>• You can restore a specific directory and the contents of that directory from a file name archive or a file system archive.<br>• You can restore an entire file system.<br>• You can restore only the permissions or only the ACL attributes of the files from the archive. |

## Managing storage using the TotalStorage Productivity Center

You can use the TotalStorage Productivity Center collect and view information about the Virtual I/O Server.

With Virtual I/O Server 1.5.2, you can install and configure the TotalStorage Productivity Center agents on the Virtual I/O Server. TotalStorage Productivity Center is an integrated, infrastructure management suite for storage that is designed to help simplify and automate the management of storage devices, storage networks, and capacity utilization of file systems and databases. When you install and configure the TotalStorage Productivity Center agents on the Virtual I/O Server, you can use the TotalStorage Productivity Center interface to collect and view information about the Virtual I/O Server. You can then perform the following tasks using the TotalStorage Productivity Center interface:

1. Run a discovery job for the agents on the Virtual I/O Server.
2. Run probes, run scans, and ping jobs to collect storage information about the Virtual I/O Server.
3. Generate reports using the Fabric Manager and the Data Manager to view the storage information gathered.
4. View the storage information gathered using the topology Viewer.

For more information, see “Installing and configuring the TotalStorage Productivity Center agents” on page 92.

# Managing networks

You can change the network configuration of the Virtual I/O Server logical partition, enable and disable GARP VLAN Registration Protocol (GVRP) on your Shared Ethernet Adapters, use Simple Network Management Protocol (SNMP) to manage systems and devices in complex networks, and upgrade to Internet Protocol version 6 (IPv6).

## Changing the network configuration of the Virtual I/O Server logical partition

Follow these steps to change or remove the network settings on the Virtual I/O Server logical partition, such as the IP address, subnet mask, gateway, and nameserver address

In this scenario, the Virtual I/O Server logical partition already has its network configuration set. The current configuration will be removed, and the updated configuration will then be set. If you plan to undo your Internet Protocol version 6 (IPv6) configuration, use the following process and commands to completely remove the TCP/IP interface and then configure a new TCP/IP interface for Internet Protocol version 4 (IPv4).

1. View the current network configuration using the lstcpip command.
2. Remove the current network configuration by running the rmtcpip command. You can remove all network settings or just the specific settings that need to be updated.
3. Configure the new network settings using the mktcpip command.

The following example is for IPv4 where the Virtual I/O Server logical partition needs to have its domain name server (DNS) information updated from its current address to 9.41.88.180:

1. Run `lstcpip -namesrv` to view the current configuration. Ensure you want to update this configuration.
2. Run `rmtcpip -namesrv` to remove the current configuration.
3. Run `mktcpip -nsrvaddr 9.41.88.180` to update the nameserver address.

## Enabling and disabling GVRP

You can enable and disable GARP VLAN Registration Protocol (GVRP) on your Shared Ethernet Adapters to control dynamic registration of VLANs over networks.

With Virtual I/O Server version 1.4, Shared Ethernet Adapters support GARP VLAN Registration Protocol (GVRP) which is based on GARP (Generic Attribute Registration Protocol). GVRP allows for the dynamic registration of VLANs over networks.

By default, GVRP is disabled on Shared Ethernet Adapters.

Before you start, create and configure the Shared Ethernet Adapter. For instructions, see “Creating a Shared Ethernet Adapter using HMC version 7” on page 84.

To enable or disable GVRP, run the following command:

```
chdev -dev Name -attr gvrp=yes/no
```

Where:

- *Name* is the name of the Shared Ethernet Adapter.
- *yes/no* defines whether GVRP is enabled or disabled. Type `yes` to enable GVRP and type `no` to disable GVRP.

## Managing SNMP on the Virtual I/O Server

Find commands for enabling, disabling, and working with SNMP on the Virtual I/O Server.

Simple Network Management Protocol (SNMP) is a set of protocols for monitoring systems and devices in complex networks. SNMP network management is based on the familiar client-server model that is widely used in Internet protocol (IP) network applications. Each managed host runs a process called an agent. The agent is a server process that maintains information about managed devices in the Management Information Base (MIB) database for the host. Hosts that are involved in network management decision-making can run a process called a manager. A manager is a client application that generates requests for MIB information and processes responses. In addition, a manager might send requests to agent servers to modify MIB information.

In general, network administrators use SNMP to more easily manage their networks for the following reasons:

- It hides the underlying system network

- The administrator can manage and monitor all network components from one console

SNMP is available on Virtual I/O Server version 1.4 and later.

The following table lists the SNMP management tasks available on the Virtual I/O Server, as well as the commands you need to run to accomplish each task.

*Table 27. Tasks and associated commands for working with SNMP on the Virtual I/O Server*

| Task | Command |
|---|---|
| Enable SNMP | startnetsvc |
| Select which SNMP agent you want to run | snmpv3_ssw |
| Issue SNMP requests to agents | cl_snmp |
| Process SNMP responses returned by agents | cl_snmp |
| Request MIB information managed by an SNMP agent | snmp_info |
| Modify MIB information managed by an SNMP agent | snmp_info |
| Generate a notification, or trap, that reports an event to the SNMP manager with a specified message | snmp_trap |
| Disable SNMP | stopnetsvc |

**Related information**

Network Management

## Upgrading the Virtual I/O Server from IPv4 to IPv6

To take advantage of enhancements, such as expanded addressing and routing simplification, use the mktcpip command to upgrade the Virtual I/O Server from Internet Protocol version 4 (IPv4) to Internet Protocol version 6 (IPv6).

IPv6 is the next generation of Internet protocol and is gradually replacing the current Internet standard, Internet Protocol version 4 (IPv4). The key IPv6 enhancement is the expansion of the IP address space from 32 bits to 128 bits, providing virtually unlimited, unique IP addresses. IPv6 provides several advantages over IPv4 including expanded routing and addressing, routing simplification, header format simplification, improved traffic control, autoconfiguration, and security.

Run the following command to upgrade from the Virtual I/O Server from IPv4 to IPv6:

```
mktcpip –auto [-interface interface]
```

where *interface* specifies which interface you want to configure for IPv6.

This command automatically performs the following tasks:

- Configures all link-local addresses for IPv6 that are currently configured for IPv4.
- Turns on the specified interfaces daemonthat support IPv6.
- Starts the ndpd-host daemon.
- Ensures that the IPv6 configuration remains intact after you reboot the Virtual I/O Server.

If you decide that you want to undo the IPv6 configuration, you must completely remove the TCP/IP interface and then configure a new TCP/IP interface for IPv4. For instructions, see “Changing the network configuration of the Virtual I/O Server logical partition” on page 100.

# Backing up the Virtual I/O Server

You can back up the Virtual I/O Server and user-defined virtual devices using the backupios command. You can also use Tivoli Storage Manager to schedule backups and store backups on another server.

The Virtual I/O Server contains the following types of information that you need to back up: the Virtual I/O Server itself and user-defined virtual devices.

- The Virtual I/O Server includes the base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata. All of this information is backed up when you use the backupios command. In situations where you plan to restore the Virtual I/O Server to the same system from which it was backed up, then backing up only the Virtual I/O Server itself is usually sufficient.
- User-defined virtual devices include metadata, like virtual devices mappings, that define the relationship between the physical environment and the virtual environment. This data can be saved to a location that is automatically backed up when you use the backupios command. In situations where you plan to restore the Virtual I/O Server to a new or different system (for example, in the event of a system failure or disaster), then you must back up both the Virtual I/O Server and user-defined virtual devices. Furthermore, in these situations, you must also back up the following components of your environment in order to fully recover your Virtual I/O Server configuration:
  - External device configurations, such as Storage Area Network (SAN) devices.
  - Resources defined on the Hardware Management Console (HMC), such as processor and memory allocations. This means backing up your HMC partition profile data for the Virtual I/O Server and its client partitions.
  - The operating systems and applications running in the client logical partitions.

You can back up and restore the Virtual I/O Server as follows.

*Table 28. Backup and restoration methods for the Virtual I/O Server*

| Backup method | Media | Restoration method |
|---|---|---|
| To tape | Tape | From tape |
| To DVD | DVD-RAM | From DVD |
| To remote file system | nim_resources.tar image | From an HMC using the Network Installation Management (NIM) on Linux facility and the installios command |
| To remote file system | mksysb image | From an AIX 5L™ NIM server and a standard mksysb system installation |
| Tivoli Storage Manager | mksysb image | Tivoli Storage Manager |

## Backing up the Virtual I/O Server to tape

You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to tape.

If the system is managed by the Integrated Virtualization Manager, then you need to back up your partition profile data for the management partition and its clients before you back up the Virtual I/O Server. For instructions, see “Backing up and restoring partition data” on page 206. (Alternatively, you can use the bkprofdata command.)

To back up the Virtual I/O Server to tape, follow these steps:

1. Assign a tape drive to the Virtual I/O Server.
2. Get the device name by typing the following command:

   ```
   lsdev -type tape
   ```

   If the tape device is in the `Defined` state, type the following command, where *dev* is the name of your tape device:

   ```
   cfgdev -dev dev
   ```

3. Type the following command, where *tape_device* is the name of the tape device you want to back up to:

```
backupios -tape tape_device
```

This command creates a bootable tape that you can use to restore the Virtual I/O Server.

4. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then you need to back up the user-defined virtual devices. For instructions, see “Backing up user-defined virtual devices” on page 106.

## Backing up the Virtual I/O Server to one or more DVDs

You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to DVD.

If the system is managed by the Integrated Virtualization Manager, then you need to back up your partition profile data for the management partition and its clients before you back up the Virtual I/O Server. For instructions, see “Backing up and restoring partition data” on page 206. (Alternatively, you can use the bkprofdata command.)

To back up the Virtual I/O Server to one or more DVDs, follow these steps. Only DVD-RAM media can be used to back up the Virtual I/O Server.

**Note:** Vendor disk drives might support burning to additional disk types, such as CD-RW and DVD-R. Refer to the documentation for your drive to determine which disk types are supported.

1. Assign an optical drive to the Virtual I/O Server logical partition.
2. Get the device name by typing the following command:

   ```
   lsdev -type optical
   ```

   If the device is in the `Defined` state, type:

   ```
   cfgdev -dev dev
   ```
3. Run the backupios command with the **-cd** option. Specify the path to the device. For example:

   ```
   backupios -cd /dev/cd0
   ```

   **Note:** If the Virtual I/O Server does not fit on one DVD, then the backupios command provides instructions for disk replacement and removal until all the volumes have been created.
   This command creates one or more bootable DVDs that you can use to restore the Virtual I/O Server.
4. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then you need to back up the user-defined virtual devices. For instructions, see “Backing up user-defined virtual devices” on page 106.

## Backing up the Virtual I/O Server to a remote file system by creating a nim_resources.tar file

You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to a remote file system by creating a nim_resources.tar file.

Backing up the Virtual I/O Server to a remote file system will create the nim_resources.tar image in the directory you specify. The nim_resources.tar file contains all the necessary resources to restore the Virtual I/O Server, including the mksysb image, the bosinst.data file, the network boot image, and Shared Product Object Tree (SPOT) resource.

The backupios command empties the target_disks_stanza section of bosinst.data and sets `RECOVER_DEVICES=Default`. This allows the mksysb file generated by the command to be cloned to another logical partition. If you plan to use the nim_resources.tar image to install to a specific disk, then you need to repopulate the target_disk_stanza section of bosinst.data and replace this file in the nim_resources.tar image. All other parts of the nim_resources.tar image must remain unchanged.

Before you start, complete the following tasks:

1. If the system is managed by the Integrated Virtualization Manager, then you need to back up your partition profile data for the management partition and its clients before you back up the Virtual I/O Server. For instructions, see “Backing up and restoring partition data” on page 206. (Alternatively, you can use the bkprofdata command.)
2. Ensure that the remote file system is available and mounted.
3. Ensure that the Virtual I/O Server has root write access to the server on which the backup will be created.

To back up the Virtual I/O Server to a remote file system, follow these steps:

1. Create a mount directory where the backup image, nim_resources.tar, will be written. For example, to create the directory /home/backup, type:
   ```
   mkdir /home/backup
   ```
2. Mount an exported directory on the mount directory. For example:
   ```
   mount server1:/export/ios_backup /home/backup
   ```
3. Run the **backupios** command with the **-file** option. Specify the path to the mounted directory. For example:
   ```
   backupios -file /home/backup
   ```

   This command creates a nim_resources.tar file that you can use to restore the Virtual I/O Server from the HMC.
4. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then you need to back up the user-defined virtual devices. For instructions, see “Backing up user-defined virtual devices” on page 106.

### Backing up the Virtual I/O Server to a remote file system by creating a mksysb image

You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to a remote file system by creating a mksysb file.

Backing up the Virtual I/O Server to a remote file system will create the mksysb image in the directory you specify. The mksysb image is an installable image of the root volume group in a file.

Before you start, complete the following tasks:

1. If the system is managed by the Integrated Virtualization Manager, then you need to back up your partition profile data for the management partition and its clients before you back up the Virtual I/O Server. For instructions, see “Backing up and restoring partition data” on page 206. (Alternatively, you can use the bkprofdata command.)
2. If you plan to restore the Virtual I/O Server from a Network Installation Management (NIM) server, verify that the NIM server is at the latest release of AIX.
3. Ensure that the remote file system is available and mounted.
4. Ensure that the Virtual I/O Server has root write access to the server on which the backup will be created.

To back up the Virtual I/O Server to a remote file system, follow these steps:

1. Create a mount directory where the backup image, mksysb image, will be written. For example, to create the directory /home/backup, type:
   ```
   mkdir /home/backup
   ```
2. Mount an exported directory on the mount directory. For example:
   ```
   mount server1:/export/ios_backup /home/backup
   ```

   where *server1* is the NIM server from which you plan to restore the Virtual I/O Server.

3. Run the backupios command with the **-file** option. Specify the path to the mounted directory. For example:

   ```
   backupios -file /home/backup/filename.mksysb -mksysb
   ```

   where *filename* is the name of mksysb image that this command creates in the specified directory. You can use the mksysb image to restore the Virtual I/O Server from a NIM server.
4. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then you need to back up the user-defined virtual devices. For instructions, see "Backing up user-defined virtual devices."

## Backing up user-defined virtual devices

In addition to backing up the Virtual I/O Server, you need to back up user-defined virtual devices (such as virtual device mappings) in preparation of a system failure or disaster.

User-defined virtual devices include metadata, such as virtual device mappings, that define the relationship between the physical environment and the virtual environment. In situations where you plan to restore the Virtual I/O Server to a new or different system (for example, in the event of a system failure or disaster), you need to back up both the Virtual I/O Server and user-defined virtual devices.

Before you start, complete the following tasks:

1. Back up the Virtual I/O Server to tape, DVD, or a remote file system. For instructions, see one of the following procedures:
   - "Backing up the Virtual I/O Server to tape" on page 103
   - "Backing up the Virtual I/O Server to one or more DVDs" on page 104
   - "Backing up the Virtual I/O Server to a remote file system by creating a nim_resources.tar file" on page 104
   - "Backing up the Virtual I/O Server to a remote file system by creating a mksysb image" on page 105
2. Decide whether you want to create a script of the following procedure. Scripting these commands makes it easy to schedule automated backups of the information.

To back up user-defined virtual devices, complete the following steps:

1. List volume groups (and storage pools) to determine what user-defined disk structures you want to back up by running the following command:

   ```
   lsvg
   ```
2. Activate each volume group (and storage pool) that you want to back up by running the following command for each volume group:

   ```
   activatevg volume_group
   ```

   where *volume_group* is the name of the volume group (or storage pool) that you want to activate.
3. Back up each volume group (and storage pool) by running the following command for each volume group:

   ```
   savevgstruct volume_group
   ```

   where *volume_group* is the name of the volume group (or storage pool) that you want to back up. This command writes a backup of the structure of a volume group (and therefore a storage pool) to the **/home/ios/vgbackups** directory.
4. Save the information about network settings, adapters, users, and security settings to the /home/padmin directory by running each command in conjunction with the tee command as follows:

   ```
   command | tee /home/padmin/filename
   ```

   Where:

- *command* is the command that produces the information you want to save.
- *filename* is the name of the file to which you want to save the information.

*Table 29. Commands that provide the information to save*

| **Command** | **Information provided (and saved)** |
|---|---|
| `cfgnamesrv -ls` | Saves all system configuration database entries related to domain name server information used by local resolver routines. |
| `entstat -all` *`devicename`*<br>*devicename* is the name of a device whose attributes or statistics you want to save. Run this command for each device whose attributes or statistics you want to save. | Saves Ethernet driver and device statistics for the device specified. |
| `hostmap -ls` | Saves all entries in the system configuration database. |
| `ioslevel` | Saves the current maintenance level of the Virtual I/O Server. |
| `lsdev -dev` *`devicename`* `-attr`<br>*devicename* is the name of a device whose attributes or statistics you want to save. Run this command for each device whose attributes or statistics you want to save. | Saves the attributes of the device specified. |
| `lsdev -type adapter` | Saves information about physical and logical adapters. |
| `lsuser` | Saves a list of all attributes of all the system users. |
| `netstat -routinfo` | Saves the routing tables, including the user-configured and current costs of each route. |
| `netstat -state` | Saves the state of all configured interfaces. |
| `optimizenet -list` | Saves characteristics of all network tuning parameters, including the current and reboot value, range, unit, type, and dependencies. |
| `viosecure -firewall view` | Saves a list of allowed ports. |
| `viosecure -view -nonint` | Saves all of the security level settings for noninteractive mode. |

## Scheduling backups of the Virtual I/O Server

You can schedule regular backups of the Virtual I/O Server and user-defined virtual devices to ensure that your backup copy accurately reflects the current configuration.

To ensure that your backup of the Virtual I/O Server accurately reflects your current running Virtual I/O Server, you should back up the Virtual I/O Server each time that its configuration changes. For example:

- Changing the Virtual I/O Server, like installing a fix pack.
- Adding, deleting, or changing the external device configuration, like changing the SAN configuration.
- Adding, deleting, or changing resource allocations and assignments for the Virtual I/O Server, like memory, processors, or virtual and physical devices.
- Adding, deleting, or changing user-defined virtual device configurations, like virtual device mappings.

Before you start, ensure that you are logged into the Virtual I/O Server as the prime administrator (padmin).

To back up the Virtual I/O Server and user-defined virtual devices, complete the following tasks:

1. Create a script for backing up the Virtual I/O Server, and save it in a directory that is accessible to the **padmin** user ID. For example, create a script called *backup* and save it in the `/home/padmin` directory. Ensure that your script includes commands for backing up the Virtual I/O Server and saving information about user-defined virtual devices.
2. Create a **crontab** file entry that runs the *backup* script on a regular interval. For example, to run *backup* every Saturday at 2:00 a.m., type the following commands:
   a. `crontab -e`
   b. `0 2 0 0 6 /home/padmin/backup`

   When you are finished, remember to save and exit.

## Backing up the Virtual I/O Server using Tivoli Storage Manager

You can use the Tivoli Storage Manager to automatically back up the Virtual I/O Server on regular intervals, or you can perform incremental backups.

### Backing up the Virtual I/O Server using Tivoli Storage Manager automated backup:

You can automate backups of the Virtual I/O Server using the crontab command and the Tivoli Storage Manager scheduler.

Before you start, complete the following tasks:

- Ensure that you configured the Tivoli Storage Manager client on the Virtual I/O Server. For instructions, see "Configuring the Tivoli Storage Manager client" on page 91.
- Ensure that you are logged into the Virtual I/O Server as the prime administrator (padmin).

To automate backups of the Virtual I/O Server, complete the following steps:

1. Write a script that creates a mksysb image of the Virtual I/O Server and save it in a directory that is accessible to the **padmin** user ID. For example, create a script called *backup* and save it in the `/home/padmin` directory. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then ensure that your script includes commands for saving information about user-defined virtual devices. For more information, see the following tasks:
   - For instructions about how to create a mksysb image, see "Backing up the Virtual I/O Server to a remote file system by creating a mksysb image" on page 105.
   - For instructions about how to save user-defined virtual devices, see "Backing up user-defined virtual devices" on page 106.
2. Create a crontab file entry that runs the *backup* script on a regular interval. For example, to create a mksysb image every Saturday at 2:00 a.m., type the following commands:
   a. `crontab -e`
   b. `0 2 0 0 6 /home/padmin/backup`

   When you are finished, remember to save and exit.
3. Work with the Tivoli Storage Manager administrator to associate the Tivoli Storage Manager client node with one or more schedules that are part of the policy domain. This task is not performed on the Tivoli Storage Manager client on the Virtual I/O Server. This task is performed by the Tivoli Storage Manager administrator on the Tivoli Storage Manager server.
4. Start the client scheduler and connect to the server schedule using the dsmc command as follows:
   ```
   dsmc -schedule
   ```
5. If you want the client scheduler to restart when the Virtual I/O Server restarts, then add the following entry to the /etc/inittab file:
   ```
   itsm::once:/usr/bin/dsmc sched > /dev/null 2>&1 # TSM scheduler
   ```

### Backing up the Virtual I/O Server using Tivoli Storage Manager incremental backup:

You can back up the Virtual I/O Server at any time by performing an incremental backup with the Tivoli Storage Manager.

Perform incremental backups in situations where the automated backup does not suit your needs. For example, before you upgrade the Virtual I/O Server, perform an incremental backup to ensure that you have a backup of the current configuration. Then, after you upgrade the Virtual I/O Server, perform another incremental backup to ensure that you have a backup of the upgraded configuration.

Before you start, complete the following tasks:

- Ensure that you configured the Tivoli Storage Manager client on the Virtual I/O Server. For instructions, see "Configuring the Tivoli Storage Manager client" on page 91.
- Ensure that you have a mksysb image of the Virtual I/O Server. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then ensure that the mksysb includes information about user-defined virtual devices. For more information, see the following tasks:
  - For instructions about how to create a mksysb image, see "Backing up the Virtual I/O Server to a remote file system by creating a mksysb image" on page 105.
  - For instructions about how to save user-defined virtual devices, see "Backing up user-defined virtual devices" on page 106.

To perform an incremental backup of the of the Virtual I/O Server, run the dsmc command. For example,

```
dsmc -incremental sourcefilespec
```

Where *sourcefilespec* is the directory path to where the mksysb file is located. For example, `/home/padmin/mksysb_image`.

## Restoring the Virtual I/O Server

You can restore the Virtual I/O Server and user-defined virtual devices using the installios command or Tivoli Storage Manager.

The Virtual I/O Server contains the following types of information that you need to restore: the Virtual I/O Server itself and user-defined virtual devices.

- The Virtual I/O Server includes the base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata. All of this information is restored when you use the installios command. In situations where you restore the Virtual I/O Server to the same system on which it was backed up, then restoring only the Virtual I/O Server itself is usually sufficient.
- User-defined virtual devices include metadata, such as virtual devices mappings, that define the relationship between the physical environment and the virtual environment. You can use this data to recreate the virtual devices. In situations where you restore the Virtual I/O Server to a new or different system (for example, in the event of a system failure or disaster), then you need to restore the Virtual I/O Server and recreate the virtual devices. Furthermore, in these situations, you also need to restore the following components of your environment in order to fully recover your Virtual I/O Server configuration:
  - External device configurations, such as Storage Area Network (SAN) devices.
  - Resources defined on the Hardware Management Console (HMC), such as processor and memory allocations. This means restoring your HMC partition profile data for the Virtual I/O Server and its client partitions.
  - The operating systems and applications running in the client logical partitions.

You can back up and restore the Virtual I/O Server as follows.

*Table 30. Backup and restoration methods for the Virtual I/O Server*

| Backup method | Media | Restoration method |
|---|---|---|
| To tape | Tape | From tape |
| To DVD | DVD-RAM | From DVD |
| To remote file system | nim_resources.tar image | From an HMC using the Network Installation Management (NIM) on Linux facility and the installios command |
| To remote file system | mksysb image | From an AIX 5L NIM server and a standard mksysb system installation |
| Tivoli Storage Manager | mksysb image | Tivoli Storage Manager |

## Restoring the Virtual I/O Server from tape

You can restore the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata from tape.

If the system is managed by the Integrated Virtualization Manager, then you need to restore your partition profile data for the management partition and its clients before you restore the Virtual I/O Server. For instructions, see “Backing up and restoring partition data” on page 206. (Alternatively, you can use the rstprofdata command.)

To restore the Virtual I/O Server from tape, follow these steps:

1. Specify the Virtual I/O Server logical partition to boot from the tape by using the bootlist command. Alternatively, you can alter the bootlist in the System Management Services (SMS).
2. Insert the tape into the tape drive.
3. From the SMS menu, select to install from the tape drive.
4. Follow the installation steps according to the system prompts.
5. If you restored the Virtual I/O Server to a different system from which it was backed up, then you need to restore the user-defined virtual devices. For instructions, see “Restoring user-defined virtual devices” on page 112.

## Restoring the Virtual I/O Server from one or more DVDs

You can restore the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata from one or more DVDs.

If the system is managed by the Integrated Virtualization Manager, then you need to restore your partition profile data for the management partition and its clients before you restore the Virtual I/O Server. For instructions, see “Backing up and restoring partition data” on page 206. (Alternatively, you can use the rstprofdata command.)

To restore the Virtual I/O Server from a one or more DVDs, follow these steps:

1. Specify the Virtual I/O Server partition to boot from the DVD by using the **bootlist** command. Alternatively, you can alter the bootlist in the System Management Services (SMS).
2. Insert the DVD into the optical drive.
3. From the SMS menu, select to install from the optical drive.
4. Follow the installation steps according to the system prompts.
5. If you restored the Virtual I/O Server to a different system from which it was backed up, then you need to restore the user-defined virtual devices. For instructions, see “Restoring user-defined virtual devices” on page 112.

## Restoring the Virtual I/O Server from the HMC using a nim_resources.tar file

You can restore the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata from a nim_resources.tar image stored in a remote file system.

If the system is managed by the Integrated Virtualization Manager, then you need to restore your partition profile data for the management partition and its clients before you restore the Virtual I/O Server. For instructions, see “Backing up and restoring partition data” on page 206. (Alternatively, you can use the rstprofdata command.)

To restore the Virtual I/O Server from a nim_resources.tar image in a file system, complete the following steps:

1. Run the installios command from the HMC command line. This restores a backup image, nim_resources.tar, that was created using the backupios command.
2. Follow the installation procedures according to the system prompts. The source of the installation images is the exported directory from the backup procedure. For example, `server1:/export/ios_backup`.
3. When the restoration is finished, open a virtual terminal connection (for example, using telnet) to the Virtual I/O Server that you restored. Some additional user input might be required.
4. If you restored the Virtual I/O Server to a different system from which it was backed up, you must restore the user-defined virtual devices. For instructions, see “Restoring user-defined virtual devices” on page 112.

## Restoring the Virtual I/O Server from a NIM server using a mksysb file

You can restore the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata from a mksysb image stored in a remote file system.

Before you start, complete the following tasks:

- Ensure that the server to which you plan to restore the Virtual I/O Server is defined as a Network Installation Management (NIM) resource.
- Ensure that the mksysb file (that contains the backup of the Virtual I/O Server) is on the NIM server.
- If the system is managed by the Integrated Virtualization Manager, then you need to restore your partition profile data for the management partition and its clients before you restore the Virtual I/O Server. For instructions, see “Backing up and restoring partition data” on page 206. (Alternatively, you can use the rstprofdata command.)

To restore the Virtual I/O Server from a mksysb image in a file system, complete the following tasks:

1. Define the mksysb file as a NIM resource, specifically, a NIM object, by running the nim command. To view a detailed description of the nim command, see nim Command. For example:

   ```
   nim -o define -t mksysb -a server=servername -alocation=/export/ios_backup/
   filename.mksysb objectname
   ```

   Where:

   - *servername* is the name of the server to which you plan to restore the Virtual I/O Server.
   - *filename* is the name of the mksysb file.
   - *objectname* is the name by which NIM registers and recognizes the mksysb file.
2. Define a Shared Product Object Tree (SPOT) resource for the mksysb file by running the nim command. For example:

   ```
   nim -o define -t spot -a server=servername -a location=/export/ios_backup/
   SPOT -a source=objectname SPOTname
   ```

   Where:

   - *servername* is the name of the server to which you plan to restore the Virtual I/O Server.

- *objectname* is the name by which NIM registers and recognizes the mksysb file.
- *SPOTname* is the name of the SPOT resource for the mksysb file.

3. Install the Virtual I/O Server from the mksysb file using the smit command. For example:

```
smit nim_bosinst
```

Ensure the following entry fields contain the following specifications.

*Table 31. Specifications for the SMIT command*

| Field | Specification |
|---|---|
| Installation TYPE | mksysb |
| SPOT | *SPOTname* from step 3 |
| MKSYSB | *objectname* from step 2 |
| Remain NIM client after install? | no |

4. Start the Virtual I/O Server logical partition. For instructions, see step 3, Boot the Virtual I/O Server, of Installing the Virtual I/O Server using NIM.
5. If you restored the Virtual I/O Server to a different system from which it was backed up, you must restore the user-defined virtual devices. For instructions, see “Restoring user-defined virtual devices.”

**Related information**

Using the NIM define operation

Defining a SPOT resource

Installing a client using NIM

## Restoring user-defined virtual devices

In addition to restoring the Virtual I/O Server, you might need to restore user-defined virtual devices (such as virtual device mappings). For example, in the event of a system failure, system migration, or disaster.

User-defined virtual devices include metadata, such as virtual device mappings, that define the relationship between the physical environment and the virtual environment. In situations where you plan to restore the Virtual I/O Server to a new or different system (for example, in the event of a system failure or disaster), you need to back up both the Virtual I/O Server and user-defined virtual devices.

Before you start, restore the Virtual I/O Server from tape, DVD, or a remote file system. For instructions, see one of the following procedures:

- “Restoring the Virtual I/O Server from tape” on page 110
- “Restoring the Virtual I/O Server from one or more DVDs” on page 110
- “Restoring the Virtual I/O Server from the HMC using a nim_resources.tar file” on page 111
- “Restoring the Virtual I/O Server from a NIM server using a mksysb file” on page 111

To restore user-defined virtual devices, complete the following steps:

1. List all of the backed-up volume groups (or storage pools) by running the following command:

```
restorevgstruct -ls
```

This command lists the files located in the **/home/ios/vgbackups** directory.

2. Run the lspv command to determine which disks are empty.
3. Restore the volume groups (or storage pools) to the empty disks by running the following command for each volume group (or storage pool):

```
restorevgstruct -vg volumegroup hdiskx
```

Where:

- *volumegroup* is the name of a volume group (or storage pool) from step 1.
- *hdiskx* is the name of an empty disk from step 2.

4. Re-create the mappings between the virtual devices and physical devices (including storage device mappings, shared Ethernet and Ethernet adapter mappings, and virtual LAN settings) using the mkvdev command. You can find mapping information in the file that you specified in the tee command from the backup procedure. For example, `/home/padmin/`*`filename`*.

## Restoring the Virtual I/O Server using Tivoli Storage Manager

You can use the Tivoli Storage Manager to restore the mksysb image of the Virtual I/O Server.

You can restore the Virtual I/O Server to the system from which it was backed up, or to a new or different system (for example, in the event of a system failure or disaster). The following procedure applies to restoring the Virtual I/O Server to the system from which it was backed up. First, you restore the mksysb image to the Virtual I/O Server using the dsmc command on the Tivoli Storage Manager client. But restoring the mksysb image does not restore the Virtual I/O Server. You then need to transfer the mksysb image to another system and convert the mksysb image to an installable format.

To restore the Virtual I/O Server to a new or different system, use one of the following procedures:

- "Restoring the Virtual I/O Server from tape" on page 110
- "Restoring the Virtual I/O Server from one or more DVDs" on page 110
- "Restoring the Virtual I/O Server from the HMC using a nim_resources.tar file" on page 111
- "Restoring the Virtual I/O Server from a NIM server using a mksysb file" on page 111

Before you start, complete the following tasks:

1. Ensure that the system to which you plan to transfer the mksysb image is running AIX.
2. Ensure that the system running AIX has a DVD-RW or CD-RW drive.
3. Ensure that AIX has the cdrecord and mkisofs RPMs downloaded and installed. To download and install the RPMs, see the AIX Toolbox for Linux Applications Web site.

**Restriction:** Interactive mode is not supported on the Virtual I/O Server. You can view session information by typing `dsmc` on the Virtual I/O Server command line.

To restore the Virtual I/O Server using Tivoli Storage Manager, complete the following tasks:

1. Determine which file you want to restore by running the dsmc command to display the files that have been backed up to the Tivoli Storage Manager server:
   ```
   dsmc -query
   ```
2. Restore the mksysb image using the dsmc command. For example:
   ```
   dsmc -restore sourcefilespec
   ```
   Where *sourcefilespec* is the directory path to the location where you want to restore the mksysb image. For example, /home/padmin/mksysb_image
3. Transfer the mksysb image to a server with a DVD-RW or CD-RW drive by running the following File Transfer Protocol (FTP) commands:
   a. Run the following command to make sure that the FTP server is started on the Virtual I/O Server: `startnetsvc ftp`
   b. Run the following command to make sure that the FTP server is started on the Virtual I/O Server: `startnetsvc ftp`
   c. Open an FTP session to the server with the DVD-RW or CD-RW drive: `ftp `*`server_hostname`*, where *server_hostname* is the hostname of the server with the DVD-RW or CD-RW drive.

   d. At the FTP prompt, change to the installation directory to the directory where you want to save the mksysb image.
   e. Set the transfer mode to binary: `binary`
   f. Turn off interactive prompting if it is on: `prompt`
   g. Transfer the mksysb image to the server: `mput mksysb_image`
   h. Close the FTP session, after transferring mksysb image, by typing `quit`.
4. Write the mksysb image to CD or DVD using the mkcd or mkdvd commands.
5. Reinstall the Virtual I/O Server using the CD or DVD that you just created. For instructions, see “Restoring the Virtual I/O Server from one or more DVDs” on page 110.

**Related reference**

mkcd Command

mkdvd Command

## Installing or replacing a PCI adapter with the system power on in Virtual I/O Server

You might need to install or replace a PCI adapter in the Virtual I/O Server logical partition or in the Integrated Virtualization Manager management partition. Use the procedure in this section to perform this task.

The Virtual I/O Server includes a PCI Hot Plug Manager that is similar to the PCI Hot Plug Manager in the AIX operating system. The PCI Hot Plug Manager allows you to hot plug PCI adapters into the server and then activate them for the logical partition without having to reboot the system. Use the PCI Hot Plug Manager for adding, identifying, or replacing PCI adapters in the system that are currently assigned to the Virtual I/O Server.

### Getting started

**Prerequisites:**

- If you are installing a new adapter, an empty system slot must be assigned to the Virtual I/O Server logical partition. This task can be done through dynamic logical partitioning (DLPAR) operations.
  - If you are using a Hardware Management Console (HMC), you must also update the logical partition profile of the Virtual I/O Server so that the new adapter is configured to the Virtual I/O Server after you restart the system.
  - If you are using the Integrated Virtualization Manager, an empty slot is probably already assigned to the Virtual I/O Server logical partition because all slots are assigned to the Virtual I/O Server by default. You only need to assign an empty slot to the Virtual I/O Server logical partition if you previously assigned all empty slots to other logical partitions.
- If you are installing a new adapter, ensure that you have the software required to support the new adapter and determine whether there are any existing PTF prerequisites to install.
- If you need help determining the PCI slot in which to place a PCI adapter, see PCI adapter placement in the system unit or expansion unit in the PCI Adapter Placement Guide.

Follow these steps to access the Virtual I/O Server, PCI Hot Plug Manager:

1. If you are using the Integrated Virtualization Manager, connect to the command-line interface.
2. Use the **diagmenu** command to open the Virtual I/O Server diagnostic menu. The menus are similar to the AIX diagnostic menus.
3. Select **Task Selection**, then press Enter.
4. At the Task Selection list, select **PCI Hot Plug Manager**.

## Installing a PCI adapter

To install a PCI adapter with the system power on in Virtual I/O Server, do the following:

1. From the PCI Hot Plug Manager, select **Add a PCI Hot Plug Adapter**, then press Enter. The Add a Hot-Plug Adapter window is displayed.
2. Select the appropriate empty PCI slot from those listed, and press Enter. A fast-blinking amber LED located at the back of the server near the adapter indicates that the slot has been identified.
3. Follow the instructions on the screen to install the adapter until the LED for the specified PCI slot is set to the Action state. The adapter installation is performed the same as in a stand-alone AIX logical partition and includes the following sequence of events:
   a. Set the adapter LED to the action state so that the indicator light for the adapter slot flashes
   b. Physically install the adapter
   c. Finish the adapter installation task in **diagmenu**.
4. Run the **cfgdev** command to configure the device for the Virtual I/O Server.

If you are installing a PCI, Fibre Channel adapter, it is now ready to be attached to a SAN and have LUNs assigned to the Virtual I/O Server for virtualization.

## Replacing a PCI Adapter

**Prerequisite:** Before you can remove or replace a storage adapter, you must unconfigure that adapter. See "Unconfiguring storage adapters" for instructions.

To replace a PCI adapter with the system power on in Virtual I/O Server, do the following:

1. From the PCI Hot Plug Manager, select **Unconfigure a Device**, then press Enter.
2. Press F4 (or Esc +4) to display the **Device Names** menu.
3. Select the adapter you are removing in the **Device Names** menu.
4. In the **Keep Definition** field, use the Tab key to answer `Yes`. In the **Unconfigure Child Devices** field, use the Tab key again to answer `YES`, then press Enter.
5. Press Enter to verify the information on the **ARE YOU SURE** screen. Successful unconfiguration is indicated by the `OK` message displayed next to the Command field at the top of the screen.
6. Press F4 (or Esc +4) twice to return to the Hot Plug Manager.
7. Select **replace/remove PCI Hot Plug adapter**.
8. Select the slot that has the device to be removed from the system.
9. Select **replace**. A fast-blinking amber LED located at the back of the machine near the adapter indicates that the slot has been identified.
10. Press Enter which places the adapter in the action state, meaning it is ready to be removed from the system.

## Unconfiguring storage adapters

Before you can remove or replace a storage adapter, you must unconfigure that adapter. Storage adapters are generally parent devices to media devices, such as disk drives or tape drives. Removing the parent requires that all attached child devices either be removed or placed in the define state.

Unconfiguring a storage adapter involves the following tasks:

- Closing all applications that are using the adapter you are removing, replacing, or moving
- Unmounting file systems
- Ensuring that all devices connected to the adapter are identified and stopped
- Listing all slots that are currently in use or a slot that is occupied by a specific adapter
- Identifying the adapter's slot location

- Making parent and child devices unavailable
- Making the adapter unavailable

If the adapter supports physical volumes that are in use by a client logical partition, then you might need to perform steps on the client logical partition before unconfiguring the storage adapter. For instructions, see “Preparing the client logical partitions.” For example, the adapter might be in use because the physical volume was used to create a virtual target device, or it might be part of a volume group used to create a virtual target device.

Follow these steps to unconfigure SCSI, SSA, and Fibre Channel storage adapters:

1. Connect to the Virtual I/O Server command-line interface.
2. Use the oem_setup_env command to close all applications that are using the adapter you are unconfiguring.
3. Type `lsslot-c pci` to list all the hot plug slots in the system unit and display their characteristics.
4. Type `lsdev -C` to list the current state of all the devices in the system unit.
5. Type `unmount` to unmount previously mounted file systems, directories, or files using this adapter.
6. Type `rmdev -l adapter -R` to make the adapter unavailable.

   **Attention:** Do not use the -d flag with the rmdev command for hot plug operations because this action removes your configuration.

## Preparing the client logical partitions

If the virtual target devices of the client logical partitions are not available, the client logical partitions can fail or they might be unable to perform I/O operations for a particular application. If you use the HMC to manage the system, you might have redundant Virtual I/O Server logical partitions, which allow for Virtual I/O Server maintenance and avoid downtime for client logical partitions. If you are replacing an adapter on the Virtual I/O Server and your client logical partition is dependent on one or more of the physical volumes accessed by that adapter, then you might need to take action on the client before you unconfigure the adapter.

The virtual target devices must be in the define state before the Virtual I/O Server adapter can be replaced. Do not remove the virtual devices permanently.

To prepare the client logical partitions so that you can unconfigure an adapter, complete the following steps depending on your situation.

*Table 32. Situations and steps for preparing the client logical partitions*

| Situation | Steps |
|---|---|
| You have redundant hardware on the Virtual I/O Server for the adapter. | No action is required on the client logical partition. |
| HMC-managed systems only: You have redundant Virtual I/O Server logical partitions that, in conjunction with virtual client adapters, provide multiple paths to the physical volume on the client logical partition. | No action is required on the client logical partition. However, path errors might be logged on the client logical partition. |
| HMC-managed systems only: You have redundant Virtual I/O Server logical partitions that, in conjunction with virtual client adapters, provide multiple physical volumes that are used to mirror a volume group. | See the procedures for your client operating system. |

Table 32. Situations and steps for preparing the client logical partitions (continued)

| Situation | Steps |
|---|---|
| You do not have redundant Virtual I/O Server logical partitions. | Shut down the client logical partition.<br>For instructions, see the following topics about shutting down logical partitions:<br>• For systems that are managed by the HMC, see “Shutting down AIX logical partitions using the HMC”, and “Shutting down Linux logical partitions using the HMC” in the *Logical Partitioning Guide*.[1]<br>• For systems that are managed by the Integrated Virtualization Manager, see “Shutting down logical partitions” on page 188. |
| **Note:**<br>1. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf . | |

## Viewing information and statistics about the Virtual I/O Server, the server, and virtual resources

You can view information and statistics about the Virtual I/O Server, the server, and virtual resources to help you manage and monitor the system, and troubleshoot problems.

The following table lists the information and statistics available on the Virtual I/O Server, as well as the commands you need to run to view the information and statistics.

Table 33. Information and associated commands for the Virtual I/O Server

| Information to view | Command |
|---|---|
| Statistics about kernel threads, virtual memory, disks, traps, and processor activity. | vmstat |
| Statistics for a Fibre Channel device driver. | fcstat |
| A summary of virtual memory usage. | svmon |
| Information about the Virtual I/O Server and the server, such as the server model, machine ID, Virtual I/O Server logical partition name and ID, and the LAN network number. | uname |

*Table 33. Information and associated commands for the Virtual I/O Server (continued)*

| **Information to view** | **Command** |
|---|---|
| Generic and device-specific statistics for an Ethernet driver or device, including the following information for a Shared Ethernet Adapter:<br>• Shared Ethernet Adapter statistics:<br>– Number of real and virtual adapters (If you are using Shared Ethernet Adapter failover, this number does not include the control channel adapter)<br>– Shared Ethernet Adapter flags<br>– VLAN IDs<br>– Information about real and virtual adapters<br>• Shared Ethernet Adapter failover statistics:<br>– High availability statistics<br>– Packet types<br>– State of the Shared Ethernet Adapter<br>– Bridging mode<br>• GARP VLAN Registration Protocol (GVRP) statistics:<br>– Bridge Protocol Data Unit (BPDU) statistics<br>– Generic Attribute Registration Protocol (GARP) statistics<br>– GARP VLAN Registration Protocol (GVRP) statistics<br>• Listing of the individual adapter statistics for the adapters associated with the Shared Ethernet Adapter | enstat |

The vmstat, fcstat, svmon, and uname commands are available with Virtual I/O Server version 1.5 or later.

# Monitoring the Virtual I/O Server

You can monitor the Virtual I/O Server using error logs or Tivoli Monitoring.

## Error logs

AIX and Linux client logical partitions log errors against failing I/O operations. Hardware errors on the client logical partitions associated with virtual devices usually have corresponding errors logged on the server. However, if the failure is within the client logical partition, there will not be errors on the server. Also, on Linux client logical partitions, if the algorithm for retrying SCSI temporary errors is different from the algorithm used by AIX, the errors might not be recorded on the server.

## Tivoli Monitoring

With Virtual I/O Server V1.3.0.1 (fix pack 8.1), you can install and configure the Tivoli Monitoring System Edition agent on the Virtual I/O Server. With Tivoli Monitoring System Edition , you can monitor the health and availability of multiple servers (including the Virtual I/O Server) from the Tivoli Enterprise Portal. Tivoli Monitoring System Edition gathers data from the Virtual I/O Server, including data about physical volumes, logical volumes, storage pools, storage mappings, network mappings, real memory, processor resources, mounted file system sizes, and so on. From the Tivoli Enterprise Portal, you can view a graphical representation of the data, use predefined thresholds to alert you on key metrics, and resolve issues based on recommendations provided by the Expert Advice feature of Tivoli Monitoring.

## Security on the Virtual I/O Server

Become familiar with the Virtual I/O Server security features.

Beginning with version 1.3 of the Virtual I/O Server, you can set security options that provide tighter security controls over your Virtual I/O Server environment. These options allow you to select a level of system security hardening and specify the settings allowable within that level. The Virtual I/O Server security feature also allows you to control network traffic by enabling the Virtual I/O Server firewall. You can configure these options using the viosecure command. To help you set up system security when you initially install the Virtual I/O Server, the Virtual I/O Server provides the configuration assistance menu. You can access the configuration assistance menu by running the cfgassist command.

Using the viosecure command, you can set, change, and view current security settings. By default, no Virtual I/O Server security levels are set; you must run the viosecure command to modify the settings.

The following sections provide an overview of these features.

### Virtual I/O Server system security hardening

The system security hardening feature protects all elements of a system by tightening security or implementing a higher level of security. Although hundreds of security configurations are possible with the Virtual I/O Server security settings, you can easily implement security controls by specifying a high, medium, or low security level.

Using the system security hardening features provided by Virtual I/O Server, you can specify values such as the following:

- Password policy settings
- usrck, pwdck, grpck, and sysck actions
- Default file-creation settings
- Settings included in the crontab command

Configuring a system at too high a security level might deny services that are needed. For example, telnet and rlogin are disabled for high level security because the login password is sent over the network unencrypted. If a system is configured at too low a security level, the system might be vulnerable to security threats. Since each enterprise has its own unique set of security requirements, the predefined High, Medium, and Low security configuration settings are best suited as a starting point for security configuration rather than an exact match for the security requirements of a particular enterprise. As you become more familiar with the security settings, you can make adjustments by choosing the hardening rules you want to apply. You can get information about the hardening rules by running the man command.

### Virtual I/O Server firewall

Using the Virtual I/O Server firewall, you can enforce limitations on IP activity in your virtual environment. With this feature, you can specify which ports and network services are allowed access to the Virtual I/O Server system. For example, if you need to restrict login activity from an unauthorized port, you can specify the port name or number and specify deny to remove it from the allow list. You can also restrict a specific IP address.

## Connecting to the Virtual I/O Server using OpenSSH

You can set up remote connections to the Virtual I/O Server using secure connections.

You can use the Open Source Secure Sockets Layer (OpenSSL) and Portable Secure Shell (OpenSSH) software to connect to the Virtual I/O Server using secure connections. For more information about OpenSSL and OpenSSH, see the OpenSSL Project and Portable SSH Web sites.

To connect to the Virtual I/O Server using OpenSSH, complete the following tasks:

1. If you are using a version of Virtual I/O Server prior to version 1.3.0, then install OpenSSH before you connect. For instructions, see "Downloading, installing, and updating OpenSSH and OpenSSL" on page 121.
2. Connect to the Virtual I/O Server. If you are using version 1.3.0 or later, then connect using either an interactive or noninteractive shell. If you are using a version prior to 1.3.0, then connect using only an interactive shell.
   - To connect using an interactive shell, type the following command from the command line of a remote system:
     ```
     ssh username@vioshostname
     ```
     where *username* is your user name for the Virtual I/O Server and *vioshostname* is the name of the Virtual I/O Server.
   - To connect using a noninteractive shell, run the following command:
     ```
     ssh username@vioshostname command
     ```
     Where:
     - *username* is your user name for the Virtual I/O Server.
     - *vioshostname* is the name of the Virtual I/O Server.
     - *command* is the command that you want to run. For example, `ioscli lsmap -all`.

       **Note:** When using a noninteractive shell, remember to use the full command form (including the `ioscli` prefix) for all Virtual I/O Server commands.
3. Authenticate SSH. If you are using version 1.3.0 or later, then authenticate using either passwords or keys. If you are using a version prior to 1.3.0, then authenticate using only passwords.
   - To authenticate using passwords, enter your user name and password when prompted by the SSH client.
   - To authenticate using keys, perform the following steps on the SSH client's operating system:
     a. Create a directory called $HOME/.ssh to store the keys. You can use RSA or DSA keys.
     b. Run the **ssh-keygen** command to generate public and private keys. For example,
        ```
        ssh-keygen -t  rsa
        ```
        This creates the following files in the $HOME/.ssh directory:
        - Private key: id_rsa
        - Public key: id_rsa.pub
     c. Run the following command to append the public key to the `authorized_keys2` file on the Virtual I/O Server:
        ```
        cat $HOME/.ssh/public_key_file | ssh username@vioshostname tee -a /home/username/.ssh/authorized_keys2
        ```
        Where:
        - *public_key_file* is the public key file that is generated in the previous step. For example, id_rsa.pub.
        - *username* is your user name for the Virtual I/O Server.
        - *vioshostname* is the name of the Virtual I/O Server.

The Virtual I/O Server might not include the latest version of OpenSSH or OpenSSL with each release. In addition, there might be OpenSSH or OpenSSL updates released in between Virtual I/O Server releases.

In these situations, you can update OpenSSH and OpenSSL on the Virtual I/O Server by downloading and installing OpenSSH and OpenSSL. For instructions, see "Downloading, installing, and updating OpenSSH and OpenSSL."

## Downloading, installing, and updating OpenSSH and OpenSSL

If you are using a Virtual I/O Server version prior to 1.3, you must download and install OpenSSH and OpenSSL software before you can connect to the Virtual I/O Server using OpenSSH. You can also use this procedure to update OpenSSH and OpenSSL on the Virtual I/O Server.

OpenSSH and OpenSSL might need to be updated on your Virtual I/O Server if the Virtual I/O Server did not include the latest version of OpenSSH or OpenSSL, or if there were OpenSSH or OpenSSL updates released in between Virtual I/O Server releases. In these situations, you can update OpenSSH and OpenSSL on the Virtual I/O Server by downloading and installing OpenSSH and OpenSSL using the following procedure.

For more information about OpenSSL and OpenSSH, see the OpenSSL Project and Portable SSH Web sites.

**Downloading the Open Source software:**

The OpenSSL software contains the encrypted library that is required to use the OpenSSH software. To download the software, complete the following tasks:

1. Download the OpenSSL RPM package to your workstation or host computer.
   a. To get the RPM package, go to the AIX Toolbox for Linux Applications Web site and click the **AIX Toolbox Cryptographic Content** link on the right side of the Web page.
   b. If you are registered to download the RPM packages, then sign in and accept the license agreement.
   c. If you are not registered to download the RPM packages, then complete the registration process and accept the license agreement. After registering, you are redirected to the download page.
   d. Select any version of the package for download: **openssl - Secure Sockets Layer and cryptography libraries and tools** and click **Download Now** to start the download.
2. Download the OpenSSH software by completing the following steps:

   **Note:** Alternatively, you can install the software from the AIX Expansion Pack.
   a. From your workstation (or host computer), go to the SourceFORGE.net Web site.
   b. Click **Download OpenSSH on AIX** to view the latest file releases.
   c. Select the appropriate download package and click **Download**.
   d. Click the openssh package (tar.Z file) to continue with the download.
3. Create a directory on the Virtual I/O Server for the Open Source software files. For example, to create an installation directory named install_ssh, run the following command: `mkdir install_ssh`.
4. Transfer the software packages to the Virtual I/O Server by running the following File Transfer Protocol (FTP) commands from the computer on which you downloaded the software packages:
   a. Run the following command to make sure that the FTP server is started on the Virtual I/O Server: `startnetsvc ftp`
   b. Open an FTP session to the Virtual I/O Server on your local host: `ftp` *`vios_server_hostname`*, where *vios_server_hostname* is the hostname of the Virtual I/O Server.
   c. At the FTP prompt, change to the installation directory to the directory that you created for the Open Source files: `cd` *`install_ssh`*, where *install_ssh* is the directory that contains the Open Source files.
   d. Set the transfer mode to binary: `binary`
   e. Turn off interactive prompting if it is on: `prompt`

f. Transfer the downloaded software to the Virtual I/O Server: `mput ssl_software_pkg`, where *ssl_software_pkg* is the software that you downloaded.
g. Close the FTP session, after transferring both software packages, by typing `quit`.

**Install the Open Source software on the Virtual I/O Server:**
To install the software, complete the following steps:

1. Run the following command from the Virtual I/O Server command line: `updateios -dev install_ssh -accept -install`, where *install_ssh* is the directory that contains the Open Source files. The installation program automatically starts the Secure Shell daemon (sshd) on the server.
2. Begin using the **ssh** and **scp** commands; no further configuration is required.

   **Restrictions:**
   - The **sftp** command is not supported on versions of Virtual I/O Server earlier than 1.3.
   - Noninteractive shells are not supported using OpenSSH with the Virtual I/O Server versions earlier than 1.3.

## Configuring Virtual I/O Server system security hardening

Set the security level to specify security hardening rules for your Virtual I/O Server system.

To implement system security hardening rules, you can use the viosecure command to specify a security level of high, medium, or low. A default set of rules is defined for each level. You can also set a level of default, which returns the system to the system standard settings and removes any level settings that have been applied.

The low level security settings are a subset of the medium level security settings, which are a subset of the high level security settings. Therefore, the *high* level is the most restrictive and provides the greatest level of control. You can apply all of the rules for a specified level or select which rules to activate for your environment. By default, no Virtual I/O Server security levels are set; you must run the viosecure command to modify the settings.

Use the following tasks to configure the system security settings.

### Setting a security level

To set a Virtual I/O Server security level of high, medium, or low, use the command `viosecure -level`. For example:

```
viosecure -level low -apply
```

### Changing the settings in a security level

To set a Virtual I/O Server security level in which you specify which hardening rules to apply for the setting, run the viosecure command interactively. For example:

1. At the Virtual I/O Server command line, type `viosecure -level high`. All the security level options (hardening rules) at that level are displayed ten at a time (pressing Enter displays the next set in the sequence).
2. Review the options displayed and make your selection by entering the numbers, separated by a comma, that you want to apply, or type **ALL** to apply all the options or **NONE** to apply none of the options.
3. Press **Enter** to display the next set of options, and continue entering your selections.

   **Note:** To exit the command without making any changes, type “q”.

### Viewing the current security setting

To display the current Virtual I/O Server security level setting use the viosecure command with the -view flag. For example:

```
viosecure -view
```

### Removing security level settings

- To unset any previously set system security levels and return the system to the standard system settings, run the following command: `viosecure -level default`
- To remove the security settings that have been applied, run the following command: `viosecure -undo`

## Configuring Virtual I/O Server firewall settings

Enable the Virtual I/O Server firewall to control IP activity.

The Virtual I/O Server firewall is not enabled by default. To enable the Virtual I/O Server firewall, you must turn it on by using the viosecure command with the -firewall option. When you enable it, the default setting is activated, which allows access for the following IP services:

- ftp
- ftp-data
- ssh
- web
- https
- rmc
- cimom

**Note:** The firewall settings are contained in the file viosecure.ctl in the /home/ios/security directory. If for some reason the viosecure.ctl file does not exist when you run the command to enable the firewall, you receive an error. You can use the -force option to enable the standard firewall default ports.

You can use the default setting or configure the firewall settings to meet the needs of your environment by specifying which ports or port services to allow. You can also turn off the firewall to deactivate the settings.

Use the following tasks at the Virtual I/O Server command line to configure the Virtual I/O Server firewall settings:

1. Enable the Virtual I/O Server firewall by running the following command:
   ```
   viosecure -firewall on
   ```
2. Specify the ports to allow or deny, by using the following command:
   ```
   viosecure -firwall allow | deny -port number
   ```
3. View the current firewall settings by running the following command:
   ```
   viosecure -firewall view
   ```
4. If you want to disable the firewall configuration, run the following command:
   ```
   viosecure -firewall off
   ```

## Configuring a Kerberos client on the Virtual I/O Server

You can configure a Kerberos client on the Virtual I/O Server to enhance security in communications across the Internet.

Before you start, ensure that the Virtual I/O Server version 1.5 or later.

Kerberos is a network authentication protocol that provides authentication for client and server applications by using a secret-key cyrptography. It negotiates authenticated, and optionally encrypted, communications between two points anywhere on the Internet. Kerberos authentication generally works as follows:

1. A Kerberos client sends a request for a ticket to the Key Distribution Center (KDC).

2. The KDC creates a ticket-granting ticket (TGT) for the client and encrypts it using the client's password as the key.
3. The KDC returns the encrypted TGT to the client.
4. The client attempts to decrypt the TGT, using its password.
5. If the client successfully decrypts the TGT (for example, if the client gives the correct password), the client keeps the decrypted TGT. The TGT indicates proof of the client's identity.

To configure a Kerberos client on the Virtual I/O Server, run the follwoing command.

```
mkkrb5clnt -c KDC_server -r realm_name \ -s Kerberos_server -d Kerberos_client
```

Where:

- *KDC_server* is the name of the KDC server.
- *realm_name* is the name of the realm to which you want to configure the Kerberos client.
- *Kerberos_server* is the fully qualified host name of the Kerberos server.
- *Kerberos_client* is the domain name of the Kerberos client.

For example:

```
mkkrb5clnt -c bob.kerberso.com -r KERBER.COM \ -s bob.kerberso.com -d testbox.com
```

In this example, you configure the Kerberos client, testbox.com, to the Kerberos server, bob.kerberso.com. The KDC is running on bob.kerberso.com.

## Managing users on the Virtual I/O Server

You can create, list, change, switch, and remove users by using Virtual I/O Server or the Tivoli Identity Manager.

When the Virtual I/O Server is installed, the only user type that is active is the prime administrator (**padmin**). The prime administrator can create additional user IDs with types of system administrator, service representative, or development engineer.

**Note:** You cannot create the prime administrator (**padmin**) user ID. It is automatically created and enabled after the Virtual I/O Server is installed.

The following table lists the user management tasks available on the Virtual I/O Server, as well as the commands you must run to accomplish each task.

*Table 34. Tasks and associated commands for working with Virtual I/O Server users*

| **Task** | **Command** |
|---|---|
| Change passwords | cfgassist |
| Create a system administrator user ID | mkuser |
| Create a service representative (SR) user ID | mkuser with the **-sr** flag |
| Create a development engineer (DE) user ID | mkuser with the **-de** flag |
| Create an LDAP user | mkuser with the **-ldap** flag |
| List a user's attributes<br>For example, determine whether a user is an LDAP user. | lsuser |
| Change a user's attributes | chuser |
| Switch to another user | su |
| Remove a user | rmuser |

You can use the Tivoli Identity Manager to automate the management of Virtual I/O Server users. Tivoli Identity Manager provides a Virtual I/O Server adapter that acts as an interface between the Virtual I/O Server and the Tivoli Identity Manager Server. The adapter acts as a trusted virtual administrator on the Virtual I/O Server, performing tasks like the following:

- Creating a user ID to authorize access to the Virtual I/O Server.
- Modifying an existing user ID to access the Virtual I/O Server.
- Removing access from a user ID. This deletes the user ID from the Virtual I/O Server.
- Suspending a user account by temporarily deactivating access to the Virtual I/O Server.
- Restoring a user account by reactivating access to the Virtual I/O Server.
- Changing a user account password on the Virtual I/O Server.
- Reconciling the user information of all current users on the Virtual I/O Server.
- Reconciling the user information of a particular user account on the Virtual I/O Server by performing a lookup.

For more information, see the Tivoli Identity Manager product manuals.

# Troubleshooting the Virtual I/O Server

Find information about diagnosing Virtual I/O Server problems and information about how to correct those problems.

This section includes information about troubleshooting the Virtual I/O Server. For information about troubleshooting the Integrated Virtualization Manager, see “Troubleshooting the Integrated Virtualization Manager” on page 206.

## Troubleshooting the Virtual I/O Server logical partition

Find information and procedures for troubleshooting and diagnosing the Virtual I/O Server logical partition.

### Troubleshooting virtual SCSI problems

Find information and procedures for troubleshooting virtual SCSI problems in the Virtual I/O Server.

For problem determination and maintenance, use the diagmenu command provided by the Virtual I/O Server.

If you are still having problems after using the diagmenu command, contact your next level of support and ask for assistance.

Refer to the AIX fast-path problem-isolation documentation in the Service provider information because, in certain cases, the diagnostic procedures described in the AIX fast-path problem-isolation documentation are not available from the diagmenu command menu.

### Correcting a failed Shared Ethernet Adapter configuration

You can troubleshoot errors that occur when you configure a Shared Ethernet Adapter, such as those that result in message 0514-040, by using the lsdev, netstat, and entstat commands.

When you configure a Shared Ethernet Adapter the configuration can fail with the following error:

```
Method error (/usr/lib/methods/cfgsea):
        0514-040 Error initializing a device into the kernel.
```

To correct the problem, complete the following steps:

1. Verify that the physical and virtual adapters that are being used to create the shared Ethernet device are available by running the following command:

   ```
   lsdev -type adapter
   ```

2. Make sure that the physical adapter is not configured. Run the following command:

   ```
   netstat -state
   ```

   The adapter must *not* show in the output.

3. Verify that the virtual adapters that are used are trunk adapters by running the following command:

   ```
   entstat -all entX | grep Trunk
   ```

4. Verify that the physical device and the virtual adapters in the Shared Ethernet Adapter are in agreement on the checksum offload setting.

   a. Determine the checksum offload setting on physical device by running the following command:

      ```
      lsdev -dev device_name -attr chksum_offload
      ```

      Where *device_name* is the name of the physical device. For example, ent0.

   b. If `chksum_offload` is set to `yes`, enable checksum offload for all of the virtual adapters in the Shared Ethernet Adapter by running the following command:

      ```
      chdev -dev device_name -attr chksum_offload=yes
      ```

      Where *device_name* is the name of a virtual adapter in the Shared Ethernet Adapter. For example, ent2.

   c. If `chksum_offload` is set to `no`, disable checksum offload for all of the virtual adapters in the Shared Ethernet Adapter by running the following command:

      ```
      chdev -dev device_name -attr chksum_offload=no
      ```

      Where *device_name* is the name of a virtual adapter in the Shared Ethernet Adapter.

   d. If there is no output, the physical device does not support checksum offload and therefore does not have the attribute. To resolve the error, disable checksum offload for all of the virtual adapters in the Shared Ethernet Adapter by running the following command:

      ```
      chdev -dev device_name -attr chksum_offload=no
      ```

      Where *device_name* is the name of a virtual adapter in the Shared Ethernet Adapter.

## Debugging problems with Ethernet connectivity

You can determine Ethernet connectivity problems by examining Ethernet statistics produced by the entstat command. Then, you can debug the problems using the starttrace and stoptrace commands.

To help debug problems with Ethernet connectivity, follow these steps:

1. Verify that the source client logical partition can ping another client logical partition on the same system without going through the Virtual I/O Server. If this fails, the problem is likely in the client logical partition's virtual Ethernet setup. If the ping is successful, proceed to the next step.
2. Start a ping on the source logical partition to a destination machine so that the packets are sent through the Virtual I/O Server. This ping will most likely fail. Proceed to the next step with the ping test running.
3. On the Virtual I/O Server, type the following command:

   ```
   entstat –all sea_adapter
   ```

   where *sea_adapter* is the name of your Shared Ethernet Adapter.

4. Verify that the VLAN ID to which the logical partition belongs is associated with the correct virtual adapter in the VLAN IDs section of the output. Examine the `ETHERNET STATISTICS` for the virtual adapter for this VLAN and verify that the packet counts under the `Receive statistics` column are increasing.

This verifies that the packets are being received by the Virtual I/O Server through the correct adapter. If the packets are not being received, the problem might be in the virtual adapter configuration. Verify the VLAN ID information for the adapters using the Hardware Management Console (HMC).

5. Examine the `ETHERNET STATISTICS` for the physical adapter for this VLAN and verify that the packet counts under the `Transmit statistics` column are increasing. This step verifies that the packets are being sent out of the Virtual I/O Server.
   - If this count is increasing, then the packets are going out of the physical adapter. Continue to step 6.
   - If this count is not increasing, then the packets are not going out of the physical adapter, and to further debug the problem, you must begin the system trace utility. Follow the instructions in step 9 to collect a system trace, statistical information, and the configuration description. Contact service and support if you need to debug the problem further.
6. Verify that the target system outside (on physical side of Virtual I/O Server) is receiving packets and sending out replies. If this is not happening, either the wrong physical adapter is associated with the Shared Ethernet Adapter or the Ethernet switch might not be configured correctly.
7. Examine the `ETHERNET STATISTICS` for the physical adapter for this VLAN and verify that the packet counts under the Receive statistics column are increasing. This step verifies that the ping replies are being received by the Virtual I/O Server. If this count is not increasing, the switch might not be configured correctly.
8. Examine the `ETHERNET STATISTICS` for the virtual adapter for this VLAN and verify that the packet counts under the `Transmit statistics` column are increasing. This step verifies that the packet is being transmitted by the Virtual I/O Server through the correct virtual adapter. If this count is not increasing, start the system trace utility. Follow the instructions in step 9 to collect a system trace, statistical information, and the configuration description. Work with service and support to debug the problem further.
9. Use the Virtual I/O Server trace utility to debug connectivity problems. Start a system trace using the starttrace command specifying the trace hook ID. The trace hook ID for Shared Ethernet Adapter is 48F. Use the stoptrace command to stop the trace. Use the cattracerpt command to read the trace log, format the trace entries, and write a report to standard output.

### Enabling noninteractive shells on Virtual I/O Server 1.3 or later

After upgrading the Virtual I/O Server to 1.3 or later, you can enable noninteractive shells using the startnetsvc command.

If you installed OpenSSH on a level of the Virtual I/O Server prior to 1.3, and then upgraded to 1.3 or later, noninteractive shells might not work because the SSH configuration file needs modification.

To enable noninteractive shells in Virtual I/O Server 1.3 or later, run the following command from the SSH client:

```
ioscli startnetsvc ssh
```

**Note:** You can run the startnetsvc command when the SSH service is running. In this situation, the command appears to fail, but is successful.

## Recovering when disks cannot be located

Learn how to recover from disks not displaying when trying to boot or install a client logical partition.

Occasionally, the disk that is needed to install the client logical partition cannot be located. In this situation, if the client is already installed, start the client logical partition. Ensure that you have the latest levels of the software and firmware. Then ensure that the **Slot number** of the virtual SCSI server adapter matches the **Remote partition virtual slot number** of the virtual SCSI client adapter.

1. Ensure that you have the latest levels of the Hardware Management Console, firmware, and Virtual I/O Server. Follow these steps:

a. To check whether you have the latest level of the HMC, see the *Installation and Configuration Guide for the Hardware Management Console*. To view the PDF file of the *Installation and Configuration Guide for the Hardware Management Console* (SA76-0084), approximately 3 MB in size, see sa76-0084.pdf .

b. Ensure that you have the latest firmware.

2. Ensure the server virtual SCSI adapter slot number is mapped correctly to the client logical partition remote slot number:

   a. In the navigation area, expand **Systems Management** → **Servers** and click the server on which the Virtual I/O Server logical partition is located.

   b. In the contents area, select the Virtual I/O Server logical partition.

   c. Click **Tasks** and select **Properties**.

   d. Click the **Virtual Adapters** tab.

   e. Click **Virtual SCSI**.

   f. If the values of the **Remote Partition** and **Remote Adapter** are **Any Partition** and **Any Partition Slot**, then complete the following steps:

      - Expand **Virtual SCSI** and click the slot number.
      - Select **Only selected client partition can connect**.
      - Enter the client logical partition's ID and adapter and click **OK**
      - Click **Virtual SCSI**.

   g. Record values of the **Remote Partition** and **Remote Adapter**. These values represent the client logical partition and the slot number of the client's virtual SCSI adapter that can connect to the associated server adapter. For example, the values of **Remote Partition**, **Remote Adapter**, and **Adapter** are as follows: AIX_client, 2, 3. This means that virtual SCSI adapter 2 on the client logical partition AIX_client can connect to the Virtual I/O Server virtual SCSI adapter 3.

   h. Repeat steps a through g for the client logical partition.

3. Ensure the server virtual SCSI adapter slot number is mapped correctly to the client logical partition remote slot number. Follow these steps:

   a. Right-click the server profile, and select **Properties**.

   b. Click the Virtual I/O Server tab.

   c. If the **Only selected remote partition and slot can connect** radio button is not selected, select it.

   d. Note the **Remote partition** and **Remote partition virtual slot number** values. This shows the client logical partition name and the client logical partition virtual slot number. This is the client logical partition and slot number that can connect to the slot given in the **Slot number** dialog box at the top of the **Virtual SCSI Adapter Properties** window.

   e. Repeat items a through e in this step for the client logical partition.

4. The **Adapter** value on the client logical partition must match the **Remote Adapter** on the Virtual I/O Server logical partition, and the **Adapter** value on the Virtual I/O Server logical partition must match the **Remote Adapter** on the client logical partition. If these numbers do not match, from the HMC, modify the profile properties to reflect the correct mapping.

5. From the Virtual I/O Server command line, type `cfgdev`.

6. Shut down and reactivate the client logical partition.

7. From the Virtual I/O Server command line, type `lsmap -all`. You see results similar to the following:

```
SVSA            Physloc                                      Client Partition ID
--------------- -------------------------------------------- ------------------
vhost0          U9113.550.10BE8DD-V1-C3                      0x00000002

VTD                   vhdisk0
LUN                   0x8100000000000000
Backing device        hdisk5
Physloc               U787B.001.DNW025F-P1-C5-T1-W5005076300C10899-L536F000000000000
```

In this example, the client logical partition ID is 2 (`0x00000002`).

**Note:** If the client logical partition is not yet installed, the Client Partition ID is 0x00000000. The slot number of the server SCSI adapter is displayed under Physloc column. The digits following the `-C` specify the slot number. In this case, the slot number is 3.

8. From the Virtual I/O Server command line, type `lsdev -virtual`. You see results similar to the following:

   ```
   name            status     description

   vhost0          Available  Virtual SCSI Server Adapter

   vhdisk0         Available  Virtual Target Device - Disk
   ```

## Troubleshooting AIX client logical partitions

Find information and procedures for troubleshooting AIX client logical partitions.

If your client partition is using virtual I/O resources, check the Service Focal Point and Virtual I/O Server first to ensure that the problem is not on the server.

On client partitions running the current level of AIX, when a hardware error is logged on the server and a corresponding error is logged on the client partition, the Virtual I/O Server provides a correlation error message in the error report.

Run the following command to gather an error report:

```
errpt -a
```

Running the **errpt** command returns results similar to the following:

```
LABEL:          VSCSI_ERR2
IDENTIFIER:     857033C6

Date/Time:       Tue Feb 15 09:18:11 2005
Sequence Number: 50
Machine Id:      00C25EEE4C00
Node Id:         vio_client53A
Class:           S
Type:            TEMP
Resource Name:   vscsi2

Description
Underlying transport error

Probable Causes
PROCESSOR

Failure Causes
PROCESSOR

        Recommended Actions
        PERFORM PROBLEM DETERMINATION PROCEDURES
Detail Data
Error Log Type
01
Reserve
00
Error Number
0006
RC
0000 0002
VSCSI Pointer
```

Compare the `LABEL`, `IDENTIFIER`, and `Error Number` values from your error report to the values in the following table to help identify the problem and determine a resolution.

*Table 35. Labels, identifiers, error numbers, problem descriptions, and resolutions of common virtual SCSI client logical partition problems*

| Label | Identifier | Error Number | Problem | Resolution |
|---|---|---|---|---|
| VSCSI_ERR2 | 857033C6 | 0006<br>RC<br>0000 0002 | The virtual SCSI server adapter on the Virtual I/O Server logical partition is not open. | Make the server adapter on the Virtual I/O Server logical partition available for use. |
| | | 001C<br>RC<br>0000 0000 | The virtual SCSI server adapter on the Virtual I/O Server logical partition has been closed abruptly. | Determine why the server adapter in the Virtual I/O Server logical partition was closed. |
| VSCSI_ERR3 | ED995F18 | 000D<br>RC<br>FFFF FFF0 | The virtual SCSI server adapter on the Virtual I/O Server logical partition is being used by another client logical partition. | Terminate the client logical partition that is using the server adapter. |
| | | 000D<br>RC<br>FFFF FFF9 | The virtual SCSI server adapter (partition number and slot number) specified in the client adapter definition does not exist. | On the HMC, correct the client adapter definition to associate it with a valid server adapter. |

## Reference information for the Virtual I/O Server

Find reference information about the Virtual I/O Server commands, configuration attributes for Tivoli agents and clients, networking statistics and attributes, and Virtual I/O Server user types.

## Virtual I/O Server and Integrated Virtualization Manager command descriptions

You can view a description of each Virtual I/O Server and Integrated Virtualization Manager command.

See the *Virtual I/O Server and Integrated Virtualization Manager Command Reference*. To view the PDF file of the *Virtual I/O Server and Integrated Virtualization Manager Command Reference* (SA76-0101), approximately 4 MB in size, see sa76-0101.pdf .

## Configuration attributes for Tivoli agents and clients

Learn about required and optional configuration attributes and variables for the Tivoli Monitoring agent, the Tivoli Usage and Accounting Manager agent, the Tivoli Storage Manager client, and the Tivoli TotalStorage Productivity Center agents.

In the following tables, the term *attribute* refers to an option that you can add to a Virtual I/O Server command. The term *variable* refers to an option that you can specify in a configuration file for Tivoli Storage Manager or Tivoli Usage and Accounting Manager.

## Tivoli Monitoring

*Table 36. Tivoli Monitoring configuration attributes*

| Attribute | Description |
|---|---|
| HOSTNAME | The host name or IP address of the Tivoli Enterprise Monitoring Server (TEMS) server to which the monitoring agent sends data. |
| MANAGING_SYSTEM | The host name or IP address of the Hardware Management Console (HMC) attached to the managed system on which the Virtual I/O Server with the monitoring agent is located. You can specify only one HMC per monitoring agent.<br>If you do not specify the MANAGING_SYSTEM attribute, the Virtual I/O Server uses the Resource Monitoring and Control (RMC) connection to obtain the host name of IP address of the HMC.<br>If the monitoring agent is running on the Integrated Virtualization Manager, then you do not need to specify the MANAGING_SYSTEM attribute. |
| RESTART_ON_REBOOT | Determines whether the monitoring agent restarts whenever the Virtual I/O Server restarts. TRUE indicates that the monitoring agent restarts whenever the Virtual I/O Server restarts. FALSE indicates that the monitoring agent does not restart whenever the Virtual I/O Server restarts. |

## Tivoli Storage Manager

*Table 37. Tivoli Storage Manager configuration attributes*

| Attribute | Description |
|---|---|
| SERVERNAME | The host name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated. |
| SERVERIP | The IP address or domain name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated. |
| NODENAME | The name of the machine on which the Tivoli Storage Manager client is installed. |

## Tivoli Usage and Accounting Manager

*Table 38. Tivoli Usage and Accounting Manager configuration variables in the A_config.par file*

| Variable | Description | Possible values | Default value |
|---|---|---|---|
| AACCT_TRANS_IDS | Designates the AIX advanced accounting record types included within the usage reports. | 1, 4, 6, 7, 8, 10, 11, or 16 | 10 |

*Table 38. Tivoli Usage and Accounting Manager configuration variables in the A_config.par file (continued)*

| Variable | Description | Possible values | Default value |
|---|---|---|---|
| AACCT_ONLY | Determines whether the Usage and Accounting Manager agent collects accounting data. | • Y: Indicates that the Usage and Accounting Manager agent collects accounting data.<br>• N: Indicates that the Usage and Accounting Manager agent does not collect accounting data. | Y |
| ITUAM_SAMPLE | Determines whether the Usage and Accounting Manager agent collects data about the storage file system. | • Y: Indicates that the Usage and Accounting Manager agent collects data about the storage file system.<br>• N: Indicates that the Usage and Accounting Manager agent does not collect data about the storage file system. | N |

*Table 39. Tivoli Usage and Accounting Manager configuration attributes*

| Attribute | Description |
|---|---|
| ACCT_DATA0 | The size, in MB, of the first data file that holds daily accounting information. |
| ACCT_DATA1 | The size, in MB, of the second data file that holds daily accounting information. |
| ISYSTEM | The time, in minutes, when the agent generates system interval records. |
| IPROCESS | The time, in minutes, when the system generates aggregate process records. |

## TotalStorage Productivity Center attributes

*Table 40. TotalStorage Productivity Center configuration attributes*

| Attribute | Description | Required or optional |
|---|---|---|
| S | Host name or IP address of the TotalStorage Productivity Center Server associated with the TotalStorage Productivity Center agent. | Required |
| A | Host name or IP address of the Agent Manager. | Required |
| devAuth | Password for authentication to the TotalStorage Productivity Center device server. | Required |
| caPass | Password for authentication to the command agent. | Required |
| caPort | Number that identifies the port for the common agent. The default is 9510. | Optional |

*Table 40. TotalStorage Productivity Center configuration attributes (continued)*

| Attribute | Description | Required or optional |
|---|---|---|
| amRegPort | Number that identifies the registration port for the Agent Manager. The default is 9511. | Optional |
| amPubPort | Number that identifies the public port for the Agent Manager. The default is 9513. | Optional |
| dataPort | Number that identifies the port for the TotalStorage Productivity Center Data server. The default is 9549. | Optional |
| devPort | Number that identifies the port of the TotalStorage Productivity Center Device server. The default is 9550. | Optional |
| newCA | The default is true. | Optional |
| oldCA | The default is false. | Optional |
| daScan | Runs a scan for the TPC_data agent after installation. The default is true. | Optional |
| daScript | Runs the script for the TPC_data agent after installation. The default is true. | Optional |
| daIntsall | Installs the TPC_data agent. The default is true. | Optional |
| faInstall | Installs the TPC_fabric agent. The default is true. | Optional |
| U | Uninstalls the TotalStorage Productivity Center agents. Possible values include:<br>• all<br>• data<br>• fabric | Optional |

**Related information**

Tivoli Monitoring 6.1 documentation

Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide

TotalStorage Productivity Center support for agents on a Virtual I/O Server

## GARP VLAN Registration Protocol statistics

Learn about Bridge Protocol Data Unit (BPDU), Generic Attribute Registration Protocol (GARP), and GARP VLAN Registration Protocol (GVRP) displayed by running the entstat -all command. You can also view examples.

BPDU refers to all protocol packets that are exchanged between the switch and the Shared Ethernet Adapter. The only bridge protocol currently available with the Shared Ethernet Adapter is GARP. GARP is a generic protocol used to exchange attribute information between two entities. The only type of GARP currently available on the Shared Ethernet Adapter is GVRP. With GVRP, the attributes exchanged are VLAN values.

## BPDU statistics

The BPDU statistics include all BPDU packets sent or received.

*Table 41. Descriptions of BPDU statistics*

| BPDU statistic | Description |
|---|---|
| Transmit | **Packets**<br>Number of packets sent.<br>**Failed packets**<br>Number of packets that could not be sent (for example, packets that could not be sent because there was no memory to allocate the outgoing packet). |
| Receive | **Packets**<br>Number of packets received.<br>**Unprocessed Packets**<br>Packets that could not be processed because the protocol was not running at the time.<br>**Non-contiguous Packets**<br>Packets that were received in several packet fragments.<br>**Packets with unknown PID**<br>Packets that had a protocol ID (PID) different than GARP. A high number is typical because the switch might be exchanging other BPDU protocol packets that the Shared Ethernet Adapter does not support.<br>**Packets with Wrong Length**<br>Packets whose specified length (in the Ethernet header) does not match the length of the Ethernet packet received. |

## GARP statistics

The GARP statistics include those BPDU packets sent or received that are of type GARP.

*Table 42. Descriptions of GARP statistics*

| GARP statistic | Description |
|---|---|
| Transmit | **Packets**<br>Number of packets sent.<br>**Failed packets**<br>Number of packets that could not be sent (for example, packets that could not be sent because there was no memory to allocate the outgoing packet).<br>**Leave All Events**<br>Packets sent with event type *Leave All*.<br>**Join Empty Events**<br>Packets sent with event type *Join Empty*<br>**Join In Events**<br>Packets sent with event type *Join In*<br>**Leave Empty Events**<br>Packets sent with event type *Leave Empty*<br>**Leave In Events**<br>Packets sent with event type *Leave In*<br>**Empty Events**<br>Packets sent with event type *Empty* |
| Receive | **Packets**<br>Number of packets received<br>**Unprocessed Packets**<br>Packets that could not be processed because the protocol was not running at the time.<br>**Packets with Unknown Attr Type:**<br>Packets with an unsupported attribute type. A high number is typical because the switch might be exchanging other GARP protocol packets that the Shared Ethernet Adapter does not support. For example, GARP Multicast Registration Protocol (GMRP).<br>**Leave All Events**<br>Packets received with event type *Leave All*<br>**Join Empty Events**<br>Packets received with event type *Join Empty*<br>**Join In Events**<br>Packets received with event type *Join In*<br>**Leave Empty Events**<br>Packets received with event type *Leave Empty*<br>**Leave In Events**<br>Packets received with event type *Leave In*<br>**Empty Events**<br>Packets received with event type *Empty* |

## GVRP statistics

The GVRP statistics include those GARP packets sent or received that are exchanging VLAN information using GVRP.

Table 43. Descriptions of GVRP statistics

| GVRP statistic | Description |
|---|---|
| Transmit | **Packets**<br>Number of packets sent<br>**Failed packets**<br>Number of packets that could not be sent (for example, packets that could not be sent because there was no memory to allocate the outgoing packet).<br>**Leave All Events**<br>Packets sent with event type *Leave All*.<br>**Join Empty Events**<br>Packets sent with event type *Join Empty*<br>**Join In Events**<br>Packets sent with event type *Join In*<br>**Leave Empty Events**<br>Packets sent with event type *Leave Empty*<br>**Leave In Events**<br>Packets sent with event type *Leave In*<br>**Empty Events**<br>Packets sent with event type *Empty* |

Table 43. Descriptions of GVRP statistics (continued)

| GVRP statistic | Description |
|---|---|
| Receive | **Packets**<br>Number of packets received.<br>**Unprocessed Packets**<br>Packets that could not be processed because the protocol was not running at the time.<br>**Packets with Invalid Length**<br>Packets that contains one or more attributes whose length does not correspond to its event type.<br>**Packets with Invalid Event**<br>Packets that contain one or more attributes whose event type is invalid.<br>**Packets with Invalid Value**<br>Packets that contain one or more attributes whose value is invalid (for example, an invalid VLAN ID).<br>**Total Invalid Attributes**<br>Sum of all of the attributes that had an invalid parameter.<br>**Total Valid Attributes**<br>Sum of all of the attributes that had no invalid parameters.<br>**Leave All Events**<br>Packets sent with event type *Leave All*.<br>**Join Empty Events**<br>Packets sent with event type *Join Empty*<br>**Join In Events**<br>Packets sent with event type *Join In*<br>**Leave Empty Events**<br>Packets sent with event type *Leave Empty*<br>**Leave In Events**<br>Packets sent with event type *Leave In*<br>**Empty Events**<br>Packets sent with event type *Empty* |

## Example statistics

Running the entstat -all command returns results similar to the following:

```
--------------------------------------------------------------
Statistics for adapters in the Shared Ethernet Adapter ent3
--------------------------------------------------------------
Number of adapters: 2
SEA Flags: 00000009
    < THREAD >
    < GVRP >
VLAN IDs :
    ent2: 1
Real Side Statistics:
    Packets received: 0
    Packets bridged: 0
    Packets consumed: 0
    Packets transmitted: 0
    Packets dropped: 0
Virtual Side Statistics:
    Packets received: 0
    Packets bridged: 0
    Packets consumed: 0
    Packets transmitted: 0
    Packets dropped: 0
Other Statistics:
    Output packets generated: 0
    Output packets dropped: 0
    Device output failures: 0
    Memory allocation failures: 0
    ICMP error packets sent: 0
    Non IP packets larger than MTU: 0
    Thread queue overflow packets: 0

--------------------------------------------------------------
Bridge Protocol Data Units (BPDU) Statistics:

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 2                                    Packets: 1370
Failed packets: 0                             Unprocessed Packets: 0
                                              Non-contiguous Packets: 0
                                              Packets w/ Unknown PID: 1370
                                              Packets w/ Wrong Length: 0

---------------------------------------------------------------
General Attribute Registration Protocol (GARP) Statistics:

Transmit Statistic:                           Receive Statistics:
-------------------                           -------------------
Packets: 2                                    Packets: 0
Failed packets: 0                             Unprocessed Packets: 0
                                              Packets w/ Unknow Attr. Type: 0

Leave All Events: 0                           Leave All Events: 0
Join Empty Events: 0                          Join Empty Events: 0
Join In Events: 2                             Join In Events: 0
Leave Empty Events: 0                         Leave Empty Events: 0
Leave In Events: 0                            Leave In Events: 0
Empty Events: 0                               Empty Events: 0

---------------------------------------------------------------
GARP VLAN Registration Protocol (GVRP) Statistics:

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 2                                    Packets: 0
Failed packets: 0                             Unprocessed Packets: 0
                                              Attributes w/ Invalid Length: 0
                                              Attributes w/ Invalid Event: 0
                                              Attributes w/ Invalid Value: 0
                                              Total Invalid Attributes: 0
                                              Total Valid Attributes: 0

Leave All Events: 0                           Leave All Events: 0
Join Empty Events: 0                          Join Empty Events: 0
Join In Events: 2                             Join In Events: 0
Leave Empty Events: 0                         Leave Empty Events: 0
Leave In Events: 0                            Leave In Events: 0
Empty Events: 0                               Empty Events: 0
```

## Network attributes

Find instructions for managing network attributes.

You can use several of the Virtual I/O Server commands, including chdev, mkvdev, and cfglnagg, to change device or network attributes. This section defines attributes that can be modified.

### Ethernet Attributes

You can modify the following Ethernet attributes.

| Attribute | Description |
|---|---|
| **Maximum Transmission Unit** (*mtu*) | Specifies maximum transmission unit (MTU). This value can be any number from 60 through 65535, but it is media dependent. |
| **Interface State** (*state*) | **detach** Removes an interface from the network interface list. If the last interface is detached, the network interface driver code is unloaded. To change the interface route of an attached interface, that interface must be detached and added again with the **chdev -dev** *Interface* **-attr** *state=detach* command.<br>**down** Marks an interface as inactive, which keeps the system from trying to transmit messages through that interface. Routes that use the interface, however, are not automatically disabled. (**chdev -dev** *Interface* **-attr** *state=down*)<br>**up** Marks an interface as active. This parameter is used automatically when setting the first address for an interface. It can also be used to enable an interface after the **chdev -dev** *Interface* **-attr** *state=up* command. |
| **Network Mask** (*netmask*) | Specifies how much of the address to reserve for subdividing networks into subnetworks.<br>The *mask* includes both the network part of the local address and the subnet part, which is taken from the host field of the address. The mask can be specified as a single hexadecimal number beginning with 0x, in standard Internet dotted-decimal notation.<br>In the 32-bit address, the mask contains bits with a value of 1 for the bit positions reserved for the network and subnet parts, and a bit with the value of 0 for the bit positions that specify the host. The mask contains the standard network portion, and the subnet segment is contiguous with the network segment. |

### Shared Ethernet Adapter attributes

You can modify the following Shared Ethernet Adapter attributes.

| Attribute | Description |
|---|---|
| **PVID** (*pvid*) | Specifies the PVID to use for the Shared Ethernet Adapter. |
| **PVID adapter** (*pvid_adapter*) | Specifies the default virtual adapter to use for non-VLAN tagged packets. |
| **Physical adapter** (*real_adapter*) | Specifies the physical adapter associated with the Shared Ethernet Adapter. |

| Attribute | Description |
|---|---|
| **Thread** (*thread*) | Activates or deactivates threading on the Shared Ethernet Adapter. Activating this option adds approximately 16 - 20% more machine cycles per transaction for MTU 1500 streaming, and approximately 31 – 38% more machine cycles per transaction for MTU 9000. The threading option adds more machine cycles per transaction at lower workloads due to the threads being started for each packet. At higher workload rates, such as full duplex or the request/response workloads, the threads can run longer without waiting and being redispatched.<br>Threaded mode should be used when virtual SCSI will be run on the same Virtual I/O Server logical partition as Shared Ethernet Adapter. Threaded mode helps ensure that virtual SCSI and the Shared Ethernet Adapter can share the processor resource appropriately. However, threading adds more instruction path length, which uses additional processor cycles. If the Virtual I/O Server logical partition will be dedicated to running shared Ethernet devices (and associated virtual Ethernet devices) only, the adapters should be configured with threading disabled.<br>You can enable or disable threading using the **-attr thread** option of the mkvdev command. To enable threading, use the `-attr thread=1` option. To disable threading, use the `-attr thread=0` option. For example, the following command disables threading for Shared Ethernet Adapter `ent1`:<br>`mkvdev -sea ent1 -vadapter ent5 -default ent5 -defaultid 1 -attr thread=0` |
| **Virtual adapters** (*virt_adapter*) | Lists the virtual Ethernet adapters associated with the Shared Ethernet Adapter. |
| **TCP segmentation offload** (*largesend*) | Enables TCP largesend capability (also known as segmentation offload) from logical partitions to the physical adapter. The physical adapter must be enabled for TCP largesend for the segmentation offload from the logical partition to the Shared Ethernet Adapter to work. Also, the logical partition must be capable of performing a largesend operation. On AIX, largesend can be enabled on a logical partition using the ifconfig command.<br>You can enable or disable TCP largesend using the -a largesend option of the chdev command. To enable it, use the ’-a largesend=1’ option. To disable it, use the ’-a largesend=0’ option.<br>For example, the following command enables *largesend* for Shared Ethernet Adapter ent1:<br>`chdev -l ent1 -a largesend=1`<br>By default the setting is disabled (largesend=0). |
| **Jumbo frames** (*jumbo_frames*) | Allows the interface configured over the Shared Ethernet Adapter to increase its MTU to 9000 bytes (the default is 1500). If the underlying physical adapter does not support jumbo frames and the *jumbo_frames* attribute is set to `yes`, then configuration fails. The underlying physical adapter must support jumbo frames. The Shared Ethernet Adapter automatically enables jumbo frames on its underlying physical adapter if *jumbo_frames* is set to `yes`. You cannot change the value of *jumbo_frames* at run time. |
| **GARP VLAN Registration Protocol (GVRP)** (*gvrp*) | Enables and disables GVRP on a Shared Ethernet Adapter. |

## Shared Ethernet Adapter failover attributes

You can modify the following Shared Ethernet Adapter failover attributes.

| Attribute | Description |
|---|---|
| **High availability mode** (*ha_mode*) | Determines whether the devices participate in a failover setup. The default is `disabled`. Typically, a Shared Ethernet Adapter in a failover setup is operating in `auto` mode, and the primary adapter is decided based on which adapter has the highest priority (lowest numerical value). A shared Ethernet device can be forced into the standby mode, where it will behave as the backup device as long as it can detect the presence of a functional primary. |
| **Control Channel** (*ctl_chan*) | Sets the virtual Ethernet device that is required for a Shared Ethernet Adapter in a failover setup so that it can communicate with the other adapter. There is no default value for this attribute, and it is required when the *ha_mode* is not set to `disabled`. |
| **Internet address to ping** (*netaddr*) | Optional attribute that can be specified for a Shared Ethernet Adapter that has been configured in a failover setup. When this attribute is specified, a shared Ethernet device will periodically ping the IP address to verify connectivity (in addition to checking for link status of the physical devices). If it detects a loss of connectivity to the specified ping host, it will initiate a failover to the backup Shared Ethernet Adapter. This attribute is not supported when you use a Shared Ethernet Adapter with a Host Ethernet Adapter (or Integrated Virtual Ethernet). |

## INET attributes

You can modify the following INET attributes.

| Attribute | Description |
|---|---|
| **Host Name** (*hostname*) | Specify the host name that you want to assign to the current machine.<br>When specifying the host name, use ASCII characters, preferably alphanumeric only. Do not use a period in the host name. Avoid using hexadecimal or decimal values as the first character (for example `3Comm`, where `3C` might be interpreted as a hexadecimal character). For compatibility with earlier hosts, use an unqualified host name of fewer than 32 characters.<br>If the host uses a domain name server for name resolution, the host name must contain the full domain name.<br>In the hierarchical domain naming system, names consist of a sequence of subnames that are not case-sensitive and that are separated by periods with no embedded blanks. The DOMAIN protocol specifies that a local domain name must be fewer than 64 characters, and that a host name must be fewer than 32 characters in length. The host name is given first. Optionally, the full domain name can be specified; the host name is followed by a period, a series of local domain names separated by periods, and finally by the root domain. A fully specified domain name for a host, including periods, must be fewer than 255 characters in length and in the following form:<br>`host.subdomain.subdomain.rootdomain`<br>In a hierarchical network, certain hosts are designated as name servers that resolve names into Internet addresses for other hosts. This arrangement has two advantages over the flat name space: resources of each host on the network are not consumed in resolving names, and the person who manages the system does not need to maintain name-resolution files on each machine on the network. The set of names managed by a single name server is known as its *zone of authority*. |
| **Gateway** (*gateway*) | Identifies the gateway to which packets are addressed. The *Gateway* parameter can be specified either by symbolic name or numeric address. |

| Attribute | Description |
|---|---|
| **Route** (*route*) | Specifies the route. The format of the *Route* attribute is: *route=destination*, *gateway*, [*metric*].<br>**destination** Identifies the host or network to which you are directing the route. The *Destination* parameter can be specified either by symbolic name or numeric address.<br>**gateway** Identifies the gateway to which packets are addressed. The *Gateway* parameter can be specified either by symbolic name or numeric address.<br>**metric** Sets the routing metric. The default is 0 (zero). The routing metric is used by the routing protocol (the *routed* daemon). Higher metrics have the effect of making a route less favorable. Metrics are counted as additional hops to the destination network or host. |

## Adapter attributes

You can modify the following adapter attributes. The attribute behavior can vary, based on the adapter and driver you have.

| Attribute | Adapters/Drivers | Description |
|---|---|---|
| **Media Speed** (*media_speed*) | • 2-Port 10/100/1000 Base-TX PCI-X Adapter<br>• 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver | The media speed attribute indicates the speed at which the adapter attempts to operate. The available speeds are 10 Mbps half-duplex, 10 Mbps full-duplex, 100 Mbps half-duplex, 100 Mbps full-duplex and autonegotiation, with a default of autonegotiation. Select auto-negotiate when the adapter should use autonegotiation across the network to determine the speed. When the network will not support autonegotiation, select the specific speed.<br>1000 MBps half and full duplex are not valid values. According to the IEEE 802.3z specification, gigabit speeds of any duplexity must be autonegotiated for copper (TX)-based adapters. If these speeds are desired, select auto-negotiate. |
| **Media Speed** (*media_speed*) | • 2-Port Gigabit Ethernet-SX PCI-X Adapter<br>• Gigabit Ethernet-SX PCI-X Adapter Device Driver | The media speed attribute indicates the speed at which the adapter attempts to operate. The available speeds are 1000 Mbps full-duplex and autonegotiation. The default is autonegotiation. Select auto-negotiate when the adapter should use autonegotiation across the network to determine the duplexity. When the network does not support autonegotiation, select 1000 Mbps full-duplex. |

| Attribute | Adapters/Drivers | Description |
|---|---|---|
| **Media Speed** (*media_speed*) | - 10/100 Mbps Ethernet PCI Adapter Device Driver | The media speed attribute indicates the speed at which the adapter attempts to operate. The available speeds are 10 Mbps half-duplex, 10 Mbps full-duplex, 100 Mbps half-duplex, 100 Mbps full-duplex and autonegotiation, with a default of autonegotiation. When the adapter should use autonegotiation across the network to determine the speed, select autonegotiate. When the network will not support autonegotiation, select the specific speed.<br>If autonegotiation is selected, the remote link device must also be set to autonegotiate to ensure the link works correctly. |
| **Media Speed** (*media_speed*) | - 10/100/1000 Base-T Ethernet PCI adapter<br>- Gigabit Ethernet-SX PCI Adapter Device Driver | The media speed attribute indicates the speed at which the adapter attempts to operate. The available speeds are 10 Mbps half-duplex, 10 Mbps full-duplex, 100 Mbps half-duplex, 100 Mbps full-duplex and autonegotiation, with a default of autonegotiation. Select autonegotiate when the adapter should use autonegotiation across the network to determine the speed. When the network will not support autonegotiation, select the specific speed.<br>For the adapter to run at 1000 Mbit/s, the autonegotiation setting must be selected.<br>**Note:** For the Gigabit Ethernet-SX PCI Adapter, the only selection available is autonegotiation. |
| **Enable Alternate Ethernet Address** (*use_alt_addr*) | | Setting this attribute to yes indicates that the address of the adapter, as it appears on the network, is the one specified by the Alternate Ethernet Address attribute. If you specify the no value, the unique adapter address written in a ROM on the adapter card is used. The default value is no. |
| **Alternate Ethernet Address** (*alt_addr*) | | Allows the adapter unique address, as it appears on the LAN network, to be changed. The value entered must be an Ethernet address of 12 hexadecimal digits and must not be the same as the address of any other Ethernet adapter. There is no default value. This field has no effect unless the Enable Alternate Ethernet Address attribute is set to yes value, in which case this field must be filled in. A typical Ethernet address is 0x02608C000001. All 12 hexadecimal digits, including leading zeros, must be entered. |
| **Enable Link Polling** (*poll_link*) | - 10/100Mbps Ethernet PCI Adapter Device Driver | Select no to cause the device driver to poll the adapter to determine the status of the link at a specified time interval. The time interval value is specified in the **Poll Link Time Interval** field. If you select no, the device driver will not poll the adapter for its link status. The default value is no. |

| Attribute | Adapters/Drivers | Description |
|---|---|---|
| **Poll Link Time Interval** (*poll_link_time*) | • 10/100Mbps Ethernet PCI Adapter Device Driver | The amount of time, in milliseconds, between polls to the adapter for its link status that the device driver is allowed. This value is required when the **Enable Link Polling** option is set to `yes`. A value between 100 through 1000 can be specified. The incremental value is 10. The default value is 500. |
| **Flow Control** (*flow_ctrl*) | • 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver<br>• Gigabit Ethernet-SX PCI-X Adapter Device Driver<br>• 2-Port 10/100/1000 Base-TX PCI-X Adapter<br>• 2-Port Gigabit Ethernet-SX PCI-X Adapter<br>• Gigabit Ethernet-SX PCI Adapter Device Driver | This attribute specifies whether the adapter should enable transmit and receive flow control. The default value is `no`. |
| **Transmit Jumbo Frames** (*jumbo_frames*) | • 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver<br>• Gigabit Ethernet-SX PCI-X Adapter Device Driver<br>• 2-Port 10/100/1000 Base-TX PCI-X Adapter<br>• 2-Port Gigabit Ethernet-SX PCI-X Adapter<br>• Gigabit Ethernet-SX PCI Adapter Device Driver | Setting this attribute to `yes` indicates that frames up to 9018 bytes in length might be transmitted on this adapter. If you specify `no`, the maximum size of frames transmitted is 1518 bytes. Frames up to 9018 bytes in length can always be received on this adapter. |

| Attribute | Adapters/Drivers | Description |
|---|---|---|
| **Checksum Offload** (*chksum_offload*) | • 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver<br>• Gigabit Ethernet-SX PCI-X Adapter Device Driver<br>• 2-Port 10/100/1000 Base-TX PCI-X Adapter<br>• 2-Port Gigabit Ethernet-SX PCI-X Adapter<br>• Gigabit Ethernet-SX PCI Adapter Device Driver<br>• Virtual Ethernet adapters | Setting this attribute to `yes` indicates that the adapter calculates the checksum for transmitted and received TCP frames. If you specify `no`, the checksum will be calculated by the appropriate software.<br>When a virtual Ethernet adapter has checksum offload enabled, the adapter advertises it to the hypervisor. The hypervisor tracks which virtual Ethernet adapters have checksum offload enabled and manages inter-partition communication accordingly.<br>When network packets are routed through the Shared Ethernet Adapter, there is a potential for link errors. In this environment, the packets must traverse the physical link with a checksum. Communication works in the following way:<br>• When a packet is received from the physical link, the physical adapter verifies the checksum. If the packet's destination is a virtual Ethernet adapter with checksum offload enabled, the receiver does not have to perform checksum verification. A receiver that does not have checksum offload enabled will accept the packet after checksum verification.<br>• When a packet originates from a virtual Ethernet adapter with checksum offload enabled, it travels to the physical adapter without a checksum. The physical adapter will generate a checksum before sending the packet out. Packets originating from a virtual Ethernet adapter with checksum offload disabled generate the checksum at the source.<br>To enable checksum offload for a Shared Ethernet Adapter, all constituent devices must have it enabled as well. The shared Ethernet device will fail if the underlying devices do not have the same checksum offload settings. |
| **Enable Hardware Transmit TCP Resegmentation** (*large_send*) | • 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver<br>• Gigabit Ethernet-SX PCI-X Adapter Device Driver<br>• 2-Port 10/100/1000 Base-TX PCI-X Adapter<br>• 2-Port Gigabit Ethernet-SX PCI-X Adapter<br>• Gigabit Ethernet-SX PCI Adapter Device Driver | This attribute specifies whether the adapter is to perform transmit TCP resegmentation for TCP segments. The default value is `no`. |

## Link Aggregation (EtherChannel) device attributes

You can modify the following Link Aggregation, or EtherChannel, attributes.

| Attribute | Description |
|---|---|
| **Link Aggregation adapters** (*adapter_names*) | The adapters that currently make up the Link Aggregation device. If you want to modify these adapters, modify this attribute and select all the adapters that should belong to the Link Aggregation device. When you use this attribute to select all of the adapters that should belong to the Link Aggregation device, its interface must not have an IP address configured. |
| **Mode** (*mode*) | The type of channel that is configured. In standard mode, the channel sends the packets to the adapter based on an algorithm (the value used for this calculation is determined by the Hash Mode attribute). In round_robin mode, the channel gives one packet to each adapter before repeating the loop. The default mode is standard.<br>Using the 802.3ad mode, the Link Aggregation Control Protocol (LACP) negotiates the adapters in the Link Aggregation device with an LACP-enabled switch.<br>If the Hash Mode attribute is set to anything other than the default, this attribute must be set to standard or 802.3ad. Otherwise, the configuration of the Link Aggregation device will fail. |
| **Hash Mode** (*hash_mode*) | If operating under standard or IEEE 802.3ad mode, the hash mode attribute determines how the outgoing adapter for each packet is chosen. Following are the different modes:<br>• `default`: uses the destination IP address to determine the outgoing adapter.<br>• `src_port`: uses the source TCP or UDP port for that connection.<br>• `dst_port`: uses the destination TCP or UDP port for that connection.<br>• `src_dst_port`: uses both the source and destination TCP or UDP ports for that connection to determine the outgoing adapter.<br>You cannot use round-robin mode with any hash mode value other than default. The Link Aggregation device configuration will fail if you attempt this combination.<br>If the packet is not TCP or UDP, it uses the default hashing mode (destination IP address).<br>Using TCP or UDP ports for hashing can make better use of the adapters in the Link Aggregationdevice, because connections to the same destination IP address can be sent over different adapters (while still retaining the order of the packets), thus increasing the bandwidth of the Link Aggregation device. |
| **Internet Address to Ping** (*netaddr*) | This field is optional. The IP address that the Link Aggregation device should ping to verify that the network is up. This is only valid when there is a backup adapter and when there are one or more adapters in the Link Aggregation device. An address of zero (or all zeros) is ignored and disables the sending of ping packets if a valid address was previously defined. The default is to leave this field blank. |
| **Retry Timeout** (*retry_time*) | This field is optional. It controls how often the Link Aggregation device sends out a ping packet to poll the current adapter for link status. This is valid only when the Link Aggregation device has one or more adapters, a backup adapter is defined, and the **Internet Address to Ping** field contains a non-zero address. Specify the timeout value in seconds. The range of valid values is 1 to 100 seconds. The default value is 1 second. |
| **Number of Retries** (*num_retries*) | This field is optional. It specifies the number of lost ping packets before the Link Aggregation device switches adapters. This is valid only when the Link Aggregation device has one or more adapters, a backup adapter is defined, and the **Internet Address to Ping** field contains a non-zero address. The range of valid values is 2 to 100 retries. The default value is 3. |

| Attribute | Description |
|---|---|
| **Enable Gigabit Ethernet Jumbo Frames** (*use_jumbo_frame*) | This field is optional. To use this attribute, all of the underlying adapters, as well as the switch, must support jumbo frames. This will work only with a Standard Ethernet (en) interface, not an IEEE 802.3 (et) interface. |
| **Enable Alternate Address** (*use_alt_addr*) | This field is optional. If you set this to yes, you can specify a MAC address that you want the Link Aggregation device to use. If you set this option to no, the Link Aggregation device uses the MAC address of the first adapter. |
| **Alternate Address** (*alt_addr*) | If **Enable Alternate Address** is set to `yes`, specify the MAC address that you want to use. The address you specify must start with `0x` and be a 12-digit hexadecimal address. |

### VLAN attributes

You can modify the following VLAN attributes.

| Attribute | Value |
|---|---|
| **VLAN Tag ID** (*vlan_tag_id*) | The unique ID associated with the VLAN driver. You can specify from 1 to 4094. |
| **Base Adapter** (*base_adapter*) | The network adapter to which the VLAN device driver is connected. |

## Shared Ethernet Adapter failover statistics

Learn about Shared Ethernet Adapter failover statistics, such as high availability information and packet types, and view examples.

### Statistic descriptions

*Table 44. Descriptions of Shared Ethernet Adapter failover statistics*

| Statistic | Description |
|---|---|
| High availability | **Control Channel PVID**<br>Port VLAN ID of the virtual Ethernet adapter used as the control channel.<br>**Control Packets in**<br>Number of packets received on the control channel.<br>**Control Packets out**<br>Number of packets sent on the control channel. |

*Table 44. Descriptions of Shared Ethernet Adapter failover statistics (continued)*

| Statistic | Description |
|---|---|
| Packet types | **Keep-Alive Packets**<br>Number of keep-alive packets received on the control channel. Keep-alive packets are received on the backup Shared Ethernet Adapter while the primary Shared Ethernet Adapter is active.<br>**Recovery Packets**<br>Number of recovery packets received on the control channel. Recovery packets are sent by the primary Shared Ethernet Adapter when it recovers from a failure and is ready to be active again.<br>**Notify Packets**<br>Number of notify packets received on the control channel. Notify packets are sent by the backup Shared Ethernet Adapter when it detects that the primary Shared Ethernet Adapter has recovered.<br>**Limbo Packets**<br>Number of limbo packets received on the control channel. Limbo packets are sent by the primary Shared Ethernet Adapter when it detects that its physical network is not operational, or when it cannot ping the specified remote host (to inform the backup that it needs to become active). |

Table 44. Descriptions of Shared Ethernet Adapter failover statistics (continued)

| Statistic | Description |
|---|---|
| State | The current state of the Shared Ethernet Adapter.<br>**INIT** The Shared Ethernet Adapter failover protocol has just been initiated.<br>**PRIMARY** The Shared Ethernet Adapter is actively connecting traffic between the VLANs to the network.<br>**BACKUP** The Shared Ethernet Adapter is idle and not connecting traffic between the VLANs and the network.<br>**RECOVERY** The primary Shared Ethernet Adapter recovered from a failure and is ready to be active again.<br>**NOTIFY** The backup Shared Ethernet Adapter detected that the primary Shared Ethernet Adapter recovered from a failure and that it needs to become idle again.<br>**LIMBO** One of the following situations is true:<br>• The physical network is not operational.<br>• The physical network's state is unknown.<br>• The Shared Ethernet Adapter cannot ping the specified remote host. |
| Bridge Mode | Describes to what level, if any, the Shared Ethernet Adapter is currently bridging traffic.<br>**Unicast** The Shared Ethernet Adapter is only sending and receiving unicast traffic (no multicast or broadcast traffic). To avoid broadcast storms, the Shared Ethernet Adapter sends and receives unicast traffic only while it is in the INIT or the RECOVERY states.<br>**All** The Shared Ethernet Adapter is sending and receiving all types of network traffic.<br>**None** The Shared Ethernet Adapter is not sending or receiving any network traffic. |
| Number of Times Server became Backup | Number of times the Shared Ethernet Adapter was active and became idle because of a failure. |
| Number of Times Server became Primary | Number of times the Shared Ethernet Adapter was idle and became active because the primary Shared Ethernet Adapter failed. |

Table 44. Descriptions of Shared Ethernet Adapter failover statistics (continued)

| Statistic | Description |
|---|---|
| High Availability Mode | How the Shared Ethernet Adapter behaves regarding the Shared Ethernet Adapter failover protocol.<br>**Auto** The Shared Ethernet Adapter failover protocol determines whether the Shared Ethernet Adapter acts as the primary Shared Ethernet Adapter or as the backup Shared Ethernet Adapter.<br>**Standby**<br>The Shared Ethernet Adapter operates as a backup if there is another Shared Ethernet Adapter available to act as the primary. *Standby* causes a primary Shared Ethernet Adapter to become a backup Shared Ethernet Adapter if there is another Shared Ethernet Adapter that can become the primary Shared Ethernet Adapter.<br>**Priority**<br>Specifies the trunk priority of the virtual Ethernet adapters of the Shared Ethernet Adapter. It is used by the Shared Ethernet Adapter protocol to determine which Shared Ethernet Adapter acts as the primary Shared Ethernet Adapter and which Shared Ethernet Adapter acts as the backup Shared Ethernet Adapter. Values range from 1 to 12, where a lower number is favored to act as a primary Shared Ethernet Adapter. |

## Example statistics

Running the entstat -all command returns results similar to the following:

```
ETHERNET STATISTICS (ent8) :
Device Type: Shared Ethernet Adapter
Hardware Address: 00:0d:60:0c:05:00
Elapsed Time: 3 days 20 hours 34 minutes 26 seconds

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 7978002                              Packets: 5701362
Bytes: 919151749                              Bytes: 664049607
Interrupts: 3                                 Interrupts: 5523380
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0
Max Packets on S/W Transmit Queue: 2
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 1

Elapsed Time: 0 days 0 hours 0 minutes 0 seconds
Broadcast Packets: 5312086                    Broadcast Packets: 3740225
Multicast Packets: 265589                     Multicast Packets: 194986
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
```

```
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 1

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 0
Driver Flags: Up Broadcast Running
 Simplex 64BitSupport ChecksumOffLoad
  DataRateSet

--------------------------------------------------------------
Statistics for adapters in the Shared Ethernet Adapter ent8
--------------------------------------------------------------
Number of adapters: 2
SEA Flags: 00000001
    < THREAD >
VLAN IDs :
    ent7: 1
Real Side Statistics:
    Packets received: 5701344
    Packets bridged: 5673198
    Packets consumed: 3963314
    Packets fragmented: 0
    Packets transmitted: 28685
    Packets dropped: 0
Virtual Side Statistics:
    Packets received: 0
    Packets bridged: 0
    Packets consumed: 0
    Packets fragmented: 0
    Packets transmitted: 5673253
    Packets dropped: 0
Other Statistics:
    Output packets generated: 28685
    Output packets dropped: 0
    Device output failures: 0
    Memory allocation failures: 0
    ICMP error packets sent: 0
    Non IP packets larger than MTU: 0
    Thread queue overflow packets: 0
High Availability Statistics:
    Control Channel PVID: 99
    Control Packets in: 0
    Control Packets out: 818825
Type of Packets Received:
    Keep-Alive Packets: 0
    Recovery Packets: 0
    Notify Packets: 0
    Limbo Packets: 0
    State: LIMBO
    Bridge Mode: All
    Number of Times Server became Backup: 0
    Number of Times Server became Primary: 0
    High Availability Mode: Auto
    Priority: 1

--------------------------------------------------------------
Real Adapter: ent2

ETHERNET STATISTICS (ent2) :
Device Type: 10/100 Mbps Ethernet PCI Adapter II (1410ff01)
Hardware Address: 00:0d:60:0c:05:00
```

```
Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 28684                                Packets: 5701362
Bytes: 3704108                                Bytes: 664049607
Interrupts: 3                                 Interrupts: 5523380
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0

Max Packets on S/W Transmit Queue: 2
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 1

Broadcast Packets: 21                         Broadcast Packets: 3740225
Multicast Packets: 0                          Multicast Packets: 194986
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 1

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 200
Driver Flags: Up Broadcast Running
 Simplex Promiscuous AlternateAddress
 64BitSupport ChecksumOffload PrivateSegment LargeSend DataRateSet

10/100 Mbps Ethernet PCI Adapter II (1410ff01) Specific Statistics:
------------------------------------------------------------------------
Link Status: Up
Media Speed Selected: Auto negotiation
Media Speed Running: 100 Mbps Full Duplex
Receive Pool Buffer Size: 1024
No Receive Pool Buffer Errors: 0
Receive Buffer Too Small Errors: 0
Entries to transmit timeout routine: 0
Transmit IPsec packets: 0
Transmit IPsec packets dropped: 0
Receive IPsec packets: 0
Receive IPsec SA offload count: 0
Transmit Large Send packets: 0
Transmit Large Send packets dropped: 0
Packets with Transmit collisions:
  1 collisions: 0     6 collisions: 0     11 collisions: 0
  2 collisions: 0     7 collisions: 0     12 collisions: 0
  3 collisions: 0     8 collisions: 0     13 collisions: 0
  4 collisions: 0     9 collisions: 0     14 collisions: 0
  5 collisions: 0    10 collisions: 0     15 collisions: 0

--------------------------------------------------------------
Virtual Adapter: ent7

ETHERNET STATISTICS (ent7) :
Device Type: Virtual I/O Ethernet Adapter (l-lan)
Hardware Address: 8a:83:54:5b:4e:9a

Transmit Statistics:                          Receive Statistics:
```

```
--------------------                          -------------------
Packets: 7949318                              Packets: 0
Bytes: 915447641                              Bytes: 0
Interrupts: 0                                 Interrupts: 0
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0

Max Packets on S/W Transmit Queue: 0
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Broadcast Packets: 5312065                    Broadcast Packets: 0
Multicast Packets: 265589                     Multicast Packets: 0
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 20000
Driver Flags: Up Broadcast Running
 Simplex Promiscuous AllMulticast
 64BitSupport ChecksumOffload DataRateSet

Virtual I/O Ethernet Adapter (l-lan) Specific Statistics:
---------------------------------------------------------
RQ Lingth: 4481
No Copy Buffers: 0
Trunk Adapter: True
  Priority: 1  Active: True
Filter MCast Mode: False
Filters: 255
  Enabled: 1  Queued: 0  Overflow: 0
LAN State: Operational

Hypervisor Send Failures: 2371664
  Receiver Failures: 2371664
  Send Errors: 0

Hypervisor Receive Failures: 0

ILLAN Attributes: 0000000000003103 [0000000000003103]

PVID:  1      VIDs: None

Switch ID: ETHERNET0

Buffers   Reg  Alloc  Min  Max   MaxA  LowReg
 tiny     512  512    512  2048  512   512
 small    512  512    512  2048  512   512
 medium   128  128    128  256   128   128
 large    24   24     24   64    24    24
 huge     24   24     24   64    24    24

--------------------------------------------------------------
Control Adapter: ent9
```

```
ETHERNET STATISTICS (ent9) :
Device Type: Virtual I/O Ethernet Adapter (l-lan)
Hardware Address: 8a:83:54:5b:4e:9b

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 821297                               Packets: 0
Bytes: 21353722                               Bytes: 0
Interrupts: 0                                 Interrupts: 0
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0

Max Packets on S/W Transmit Queue: 0
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Broadcast Packets: 821297                     Broadcast Packets: 0
Multicast Packets: 0                          Multicast Packets: 0
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 20000
Driver Flags: Up Broadcast Running
   Simplex 64BitSupport ChecksumOffload DataRateSet

Virtual I/O Ethernet Adapter (l-lan) Specific Statistics:
----------------------------------------------------------
RQ Length: 4481
No Copy Buffers: 0
Trunk Adapter: False
Filter MCast Mode: False
Filters: 255
  Enabled: 0  Queued: 0  Overflow: 0
LAN State: Operational

Hypervisor Send Failures: 0
  Receiver Failures: 0
  Send Errors: 0

Hypervisor Receive Failures: 0

ILLAN Attributes: 0000000000003002 [0000000000003002]

PVID:  99     VIDs:  None

Switch ID: ETHERNET0

Buffers        Reg  Alloc    Min     Max   MaxA  LowReg
 tiny          512    512    512    2048    512     512
 small         512    512    512    2048    512     512
```

```
medium        128    128    128     256    128     128
large          24     24     24      64     24      24
huge           24     24     24      64     24      24
```

## Shared Ethernet Adapter statistics

Learn about general Shared Ethernet Adapter statistics, such as VLAN IDs and packet information, and view examples.

### Statistic descriptions

*Table 45. Descriptions of Shared Ethernet Adapter statistics*

| Statistic | Description |
|---|---|
| Number of adapters | Includes the real adapter and all of the virtual adapters.<br>**Note:** If you are using Shared Ethernet Adapter failover, then the control channel adapter is not included. |
| Shared Ethernet Adapter flags | Denotes the features that the Shared Ethernet Adapter is currently running.<br>**THREAD** The Shared Ethernet Adapter is operating in threaded mode, where incoming packets are queued and processed by different threads; its absence denotes interrupt mode, where packets are processed in the same interrupt where they are received.<br>**LARGESEND** The large send feature has been enabled on the Shared Ethernet Adapter.<br>**JUMBO_FRAMES** The jumbo frames feature has been enabled on the Shared Ethernet Adapter.<br>**GVRP** The GVRP feature has been enabled on the Shared Ethernet Adapter. |
| VLAN IDs | List of VLAN IDs that have access to the network through the Shared Ethernet Adapter (this includes PVID and all tagged VLANs). |

*Table 45. Descriptions of Shared Ethernet Adapter statistics (continued)*

| Statistic | Description |
|---|---|
| Real adapters | **Packets received**<br>Number of packets received on the physical network.<br>**Packets bridged**<br>Number of packets received on the physical network that were sent to the virtual network.<br>**Packets consumed**<br>Number of packets received on the physical network that were addressed to the interface configured over the Shared Ethernet Adapter.<br>**Packets fragmented**<br>Number of packets received on the physical network that were fragmented before being sent to the virtual network. They were fragmented because they were bigger than the outgoing adapter's Maximum Transmission Unit (MTU).<br>**Packets transmitted**<br>Number of packets sent on the physical network. This includes packets sent from the interface configured over the Shared Ethernet Adapter, as well as each packet sent from the virtual network to the physical network (including fragments).<br>**Packets dropped**<br>Number of packets received on the physical network that were dropped for one of the following reasons:<br>• The packet was the oldest packet on a thread's queue and there was no space to accommodate a newly received packet.<br>• The packet had an invalid VLAN ID and could not be processed.<br>• The packet was addressed to the Shared Ethernet Adapter interface, but its interface had no filters registered. |

*Table 45. Descriptions of Shared Ethernet Adapter statistics (continued)*

| Statistic | Description |
|---|---|
| Virtual adapters | **Packets received**<br>Number of packets received on the virtual network. In other words, the number of packets received on all of the virtual adapters.<br>**Packets bridged**<br>Number of packets received on the virtual network that were sent to the physical network.<br>**Packets consumed**<br>Number of packets received on the virtual network that were addressed to the interface configured over the Shared Ethernet Adapter.<br>**Packets fragmented**<br>Number of packets received on the virtual network that were fragmented before being sent to the physical network. They were fragmented because they were bigger than the outgoing adapter's MTU.<br>**Packets transmitted**<br>Number of packets sent on the virtual network. This includes packets sent from the interface configured over the Shared Ethernet Adapter, as well as each packet sent from the physical network to the virtual network (including fragments).<br>**Packets dropped**<br>Number of packets received on the virtual network that were dropped for one of the following reasons:<br>• The packet was the oldest packet on a thread's queue and there was no space to accommodate a newly received packet.<br>• The packet was addressed to the Shared Ethernet Adapter interface, but its interface had no filters registered. |
| Output packets generated | Number of packets with a valid VLAN tag or no VLAN tag sent out of the interface configured over the Shared Ethernet Adapter. |
| Output packets dropped | Number of packets sent out of the interface configured over the Shared Ethernet Adapter that are dropped because of an invalid VLAN tag. |
| Device output failures | Number of packets that could not be sent due to underlying device errors. This includes errors sent on the physical network and virtual network, including fragments and Internet Control Message Protocol (ICMP) error packets generated by the Shared Ethernet Adapter. |
| Memory allocation failures | Number of packets that could not be sent because there was insufficient network memory to complete an operation. |

*Table 45. Descriptions of Shared Ethernet Adapter statistics (continued)*

| Statistic | Description |
|---|---|
| ICMP error packets sent | Number of ICMP error packets successfully sent when a big packet could not be fragmented because the *don't fragment* bit was set. |
| Non IP packets larger than MTU | Number of packets that could not be sent because they were bigger than the outgoing adapter's MTU and could not be fragmented because they were not IP packets. |
| Thread queue overflow packets | Number of packets that were dropped from the thread queues because there was no space to accommodate a newly received packet. |

## Example statistics

```
ETHERNET STATISTICS (ent8) :
Device Type: Shared Ethernet Adapter
Hardware Address: 00:0d:60:0c:05:00
Elapsed Time: 3 days 20 hours 34 minutes 26 seconds

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 7978002                              Packets: 5701362
Bytes: 919151749                              Bytes: 664049607
Interrupts: 3                                 Interrupts: 5523380
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0
Max Packets on S/W Transmit Queue: 2
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 1

Elapsed Time: 0 days 0 hours 0 minutes 0 seconds
Broadcast Packets: 5312086                    Broadcast Packets: 3740225
Multicast Packets: 265589                     Multicast Packets: 194986
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 1

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 0
Driver Flags: Up Broadcast Running
 Simplex 64BitSupport ChecksumOffLoad
  DataRateSet

--------------------------------------------------------------
Statistics for adapters in the Shared Ethernet Adapter ent8
--------------------------------------------------------------
Number of adapters: 2
SEA Flags: 00000001
    < THREAD >
VLAN IDs :
    ent7: 1
```

```
Real Side Statistics:
    Packets received: 5701344
    Packets bridged: 5673198
    Packets consumed: 3963314
    Packets fragmented: 0
    Packets transmitted: 28685
    Packets dropped: 0
Virtual Side Statistics:
    Packets received: 0
    Packets bridged: 0
    Packets consumed: 0
    Packets fragmented: 0
    Packets transmitted: 5673253
    Packets dropped: 0
Other Statistics:
    Output packets generated: 28685
    Output packets dropped: 0
    Device output failures: 0
    Memory allocation failures: 0
    ICMP error packets sent: 0
    Non IP packets larger than MTU: 0
    Thread queue overflow packets: 0

--------------------------------------------------------------
Real Adapter: ent2

ETHERNET STATISTICS (ent2) :
Device Type: 10/100 Mbps Ethernet PCI Adapter II (1410ff01)
Hardware Address: 00:0d:60:0c:05:00

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 28684                                Packets: 5701362
Bytes: 3704108                                Bytes: 664049607
Interrupts: 3                                 Interrupts: 5523380
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0

Max Packets on S/W Transmit Queue: 2
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 1

Broadcast Packets: 21                         Broadcast Packets: 3740225
Multicast Packets: 0                          Multicast Packets: 194986
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 1

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 200
Driver Flags: Up Broadcast Running
 Simplex Promiscuous AlternateAddress
 64BitSupport ChecksumOffload PrivateSegment LargeSend DataRateSet

10/100 Mbps Ethernet PCI Adapter II (1410ff01) Specific Statistics:
```

```
------------------------------------------------------------------------
Link Status: Up
Media Speed Selected: Auto negotiation
Media Speed Running: 100 Mbps Full Duplex
Receive Pool Buffer Size: 1024
No Receive Pool Buffer Errors: 0
Receive Buffer Too Small Errors: 0
Entries to transmit timeout routine: 0
Transmit IPsec packets: 0
Transmit IPsec packets dropped: 0
Receive IPsec packets: 0
Receive IPsec SA offload count: 0
Transmit Large Send packets: 0
Transmit Large Send packets dropped: 0
Packets with Transmit collisions:
  1 collisions: 0     6 collisions: 0     11 collisions: 0
  2 collisions: 0     7 collisions: 0     12 collisions: 0
  3 collisions: 0     8 collisions: 0     13 collisions: 0
  4 collisions: 0     9 collisions: 0     14 collisions: 0
  5 collisions: 0    10 collisions: 0     15 collisions: 0

--------------------------------------------------------------
Virtual Adapter: ent7

ETHERNET STATISTICS (ent7) :
Device Type: Virtual I/O Ethernet Adapter (l-lan)
Hardware Address: 8a:83:54:5b:4e:9a

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 7949318                              Packets: 0
Bytes: 915447641                              Bytes: 0
Interrupts: 0                                 Interrupts: 0
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0

Max Packets on S/W Transmit Queue: 0
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Broadcast Packets: 5312065                    Broadcast Packets: 0
Multicast Packets: 265589                     Multicast Packets: 0
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 20000
Driver Flags: Up Broadcast Running
 Simplex Promiscuous AllMulticast
 64BitSupport ChecksumOffload DataRateSet

Virtual I/O Ethernet Adapter (l-lan) Specific Statistics:
---------------------------------------------------------
RQ Lingth: 4481
```

```
No Copy Buffers: 0
Trunk Adapter: True
  Priority: 1  Active: True
Filter MCast Mode: False
Filters: 255
  Enabled: 1  Queued: 0  Overflow: 0
LAN State: Operational

Hypervisor Send Failures: 2371664
  Receiver Failures: 2371664
  Send Errors: 0

Hypervisor Receive Failures: 0

ILLAN Attributes: 0000000000003103 [0000000000003103]

PVID:  1      VIDs: None

Switch ID: ETHERNET0

Buffers   Reg  Alloc  Min  Max   MaxA  LowReg
 tiny     512   512   512  2048  512   512
 small    512   512   512  2048  512   512
 medium   128   128   128  256   128   128
 large    24    24    24   64    24    24
 huge     24    24    24   64    24    24
```

## User types for the Virtual I/O Server

Learn about Virtual I/O Server user types and their user permissions.

The Virtual I/O Server has the following user types: prime administrator, system administrator, service representative user, and development engineer user. After installation, the only user type that is active is the prime administrator.

### Prime administrator

The prime administrator (**padmin**) user ID is the only user ID that is enabled after installation of the Virtual I/O Server and can run every Virtual I/O Server command. There can be only one prime administrator in the Virtual I/O Server.

### System administrator

The system administrator user ID has access to all commands except the following commands:

- lsfailedlogin
- lsgcl
- mirrorios
- mkuser
- oem_setup_env
- rmuser
- shutdown
- unmirrorios

The prime administrator can create an unlimited number of system administrator IDs.

### Service representative

Create the service representative (SR) user so that an service representative can log in to the system and perform diagnostic routines. Upon logging in, the SR user is placed directly into the diagnostic menus.

## Development engineer

Create a Development engineer (DE) user ID so that an development engineer can log in to the system and debug problems.

## View

This role is a read-only role and can perform only list-type (ls) functions. Users with this role do not have the authority to change the system configuration and do not have write permission to their home directories.

# Chapter 4. Integrated Virtualization Manager

Manage the Virtual I/O Server and client logical partitions using the Integrated Virtualization Manager.

The Integrated Virtualization Manager provides a Web-based system management interface and a command-line interface that you can use to manage some systems and some BladeCenter blade servers that use the Virtual I/O Server. On the managed system, you can create logical partitions, manage the virtual storage and virtual Ethernet, and view service information related to the server. The Integrated Virtualization Manager is included with the Virtual I/O Server, but it is activated and usable only on certain platforms, and where no Hardware Management Console (HMC) is present.

If you install the Virtual I/O Server on a supported server, and if there is no HMC attached to the server when you install the Virtual I/O Server, then the Integrated Virtualization Manager is enabled on that server. You can then use the Integrated Virtualization Manager to configure the managed system through the Virtual I/O Server.

For information about using the Virtual I/O Server on a system that is managed by the HMC, see Chapter 3, “Virtual I/O Server,” on page 13.

On POWER6 processor-based servers, you can install the following operating systems on logical partitions that are created by using the Integrated Virtualization Manager.

*Table 46. Operating system support for logical partitions on POWER6 processor-based servers managed by the Integrated Virtualization Manager*

| Operating system | POWER6 processor-based servers |
|---|---|
| AIX 5.3 or later | • 03E/4A<br>• 04E/8A |
| SUSE Linux Enterprise Server 10 Service Pack 2 or later | • JS/12<br>• JS/22 |
| SUSE Linux Enterprise Server 10 Service Pack 1 Update 1 or later | JS/22 |
| SUSE Linux Enterprise Server 10 Service Pack 1 | • 03E/4A<br>• 04E/8A |
| Red Hat Enterprise Linux version 5.2 | JS/22 |
| Red Hat Enterprise Linux version 5.1 | • 03E/4A<br>• 04E/8A<br>• JS/12<br>• JS/22 |
| Red Hat Enterprise Linux version 4.7 | JS/22 |
| Red Hat Enterprise Linux version 4.6 | • JS/12<br>• JS/22 |
| Red Hat Enterprise Linux version 4.5 | • 03E/4A<br>• 04E/8A<br>• JS/22 |

On POWER5 processor-based servers, you can install the following operating systems on logical partitions created by using the Integrated Virtualization Manager:

- AIX 5.3 (or later)
- SUSE Linux Enterprise Server 9 (or later)
- SUSE Linux Enterprise Server 10 (or later)
- Red Hat Enterprise Linux version 4 (or later)
- Red Hat Enterprise Linux version 5 (or later)

## Integrated Virtualization Manager

The *Integrated Virtualization Manager* is a browser-based system management interface for the Virtual I/O Server. The Integrated Virtualization Manager allows you to create and manage logical partitions on a single server.

The Integrated Virtualization Manager is supported only on specific server models. For a complete list of supported server models, see the *PowerVM Editions Operations Guide*. To view the PDF file of the *PowerVM Editions Operations Guide* (SA76-0100), approximately 4 MB in size, see sa76-0100.pdf .

*Virtual I/O Server* is software that provides virtual storage and shared Ethernet resources to the other logical partitions on the managed system. Virtual I/O Server is not a general purpose operating system that can run applications. Virtual I/O Server is installed on a logical partition in the place of a general purpose operating system, and is used solely to provide virtual I/O resources to other logical partitions with general purpose operating systems. You use the Integrated Virtualization Manager to specify how these resources are assigned to the other logical partitions.

To use the Integrated Virtualization Manager, you must first install Virtual I/O Server on an unpartitioned server. Virtual I/O Server automatically creates a logical partition for itself, which is called the *management partition* for the managed system. The management partition is the Virtual I/O Server logical partition that controls all of the physical I/O resources on the managed system. After you install Virtual I/O Server, you can configure a physical Ethernet adapter on the server so that you can connect to the Integrated Virtualization Manager from a computer with a Web browser.

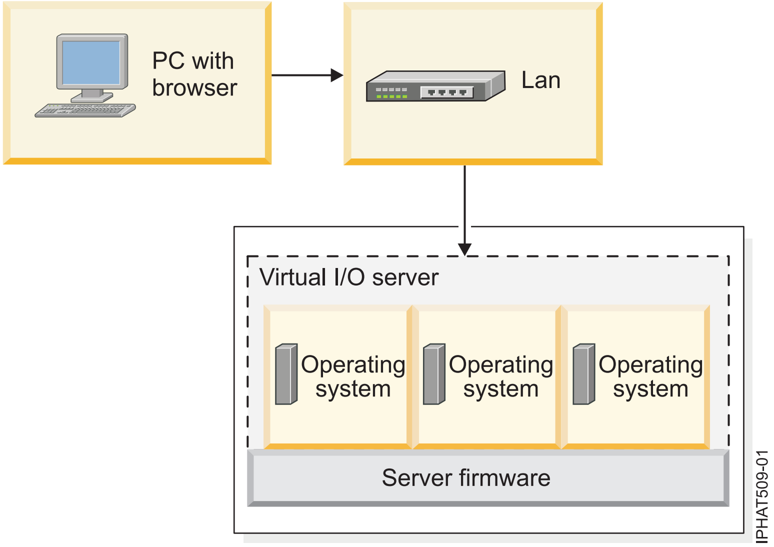

This figure illustrates an systems server or an BladeCenter blade server with POWER Architecture technology. The Virtual I/O Server is in its own logical partition, and the client logical partitions are managed by the Virtual I/O Server logical partition. The browser on the PC connects to the Integrated

Virtualization Manager interface over a network, and you can use the Integrated Virtualization Manager to create and manage the logical partitions on the server.

### Resource assignment

When you use the Integrated Virtualization Manager to create a logical partition, then you assign memory and processor resources directly to logical partitions. If you use dedicated processors, then you specify the exact number of dedicated processors. If you use shared processors, then you specify the number of virtual processors for the logical partition, and the Integrated Virtualization Manager calculates the number of processing units it assigns to the logical partition based on the number of virtual processors. In all cases, the amount of resources that you assign is committed to the logical partition from the time that you create the logical partition until the time that you change this amount or delete the logical partition. You therefore cannot overcommit processor resources to logical partitions using the Integrated Virtualization Manager.

A logical partition that is created using the Integrated Virtualization Manager has minimum and maximum memory and processor values. The minimum and maximum values are used when you use a workload management application on the managed system, when you restart the managed system after a processor failure, or when you dynamically move resources to or from the Virtual I/O Server management partition. By default, the minimum and maximum values are set to the same value as the actual amount of committed resources. You can change the minimum and maximum processor values at any time, but you can change the minimum and maximum memory values only while the logical partition is not running.

When you use the Integrated Virtualization Manager to partition your managed system, a fraction of the memory and a fraction of the processors on the managed system are assigned to the Virtual I/O Server management partition. If desired, you can change the memory and processor resources that are assigned to the management partition to match your Virtual I/O Server workload. Physical disks can be assigned directly to logical partitions, or they can be assigned to storage pools, and virtual disks (or logical volumes) can be created from these storage pools and assigned to logical partitions. Physical Ethernet connections are generally shared by configuring the physical Ethernet adapter as a virtual Ethernet bridge between the virtual LAN on the server and an external, physical LAN. Host Ethernet Adapter Other types of I/O devices

## Planning for the Integrated Virtualization Manager

You need to develop a system plan for your servers server or BladeCenter blade server that is managed by the Integrated Virtualization Manager.

Proper planning is essential for the successful setup and use of your server. When you install the Integrated Virtualization Manager, it automatically creates a logical partition for itself on the server. This logical partition is called the *management partition*. The Integrated Virtualization Manager automatically assigns a fraction of the memory and processors on the server to the management partition. You can change the default amount of memory and processor resources that are assigned to the management partition.

You need to develop a plan that includes information such as the following:

- System resource requirements for the management partition. The system resource requirements for the management partition can depend on many factors. These factors can include the server model, the number of logical partitions that you create on the managed system, and the number of virtual devices used by those logical partitions.
- Storage needs of each logical partition that you are to create on your managed system. Calculate how much storage space each logical partition requires for its operating system, applications, and data. For more information about the storage requirements for each operating system, consult the operating system documentation.

The following table provides information resources to help you create a system plan depending on your hardware model.

*Table 47. Planning information resources for servers that are managed by the Integrated Virtualization Manager*

| Hardware model | Planning information resources |
|---|---|
| POWER6 processor-based servers | • "Planning for the Virtual I/O Server" on page 45 **Note:** Although this information is focused on planning for the Virtual I/O Server on a system that is managed by a Hardware Management Console (HMC), most of the information also applies to planning for the Virtual I/O Server on a system that is managed by the Integrated Virtualization Manager. |
| BladeCenter | • Planning for PowerVM Editions (or Advanced POWER Virtualization Editions) in the Systems Information Center<br>• "Planning for the Virtual I/O Server" on page 45 **Note:** Although this information is focused on planning for the Virtual I/O Server on an servers server that is managed by an HMC, most of the information also applies to planning for the Virtual I/O Server on an BladeCenter blade server that is managed by the Integrated Virtualization Manager.<br>• The *Virtual I/O Server and Integrated Virtualization Manager Command Reference* contains a detailed description of the mksysplan command on the Integrated Virtualization Manager.<br>After you have set up and configured a system, you can use the mksysplan command to create a system plan based on the existing system configuration. Then, you can export the system plan and import it into another system managed by the Integrated Virtualization Manager. Finally, you can use the Deploy System Plan wizard on the Integrated Virtualization Manager to deploy the system plan to the new system. |

## System plan validation on the Integrated Virtualization Manager

You deploy a system plan on an Integrated Virtualization Manager managed system by using the System Plan Deployment Wizard. The wizard validates the information in the system plan against the configuration of the managed system before beginning the deployment process.

The validation process for a system plan consists of two validation phases: the hardware validation phase and the partition validation phase.

When validating the hardware on the managed system, the Integrated Virtualization Manager compares the following information from the system plan with the hardware available on the managed system:

- Processor, memory, and storage are available on the managed system
- Physical I/O adapter placement

The hardware described in the system plan passes validation if it matches the hardware specified by the managed system. The hardware on the managed system can contain resources in addition to those specified in the system plan and still pass validation, but the hardware on the managed system must at least match the hardware specified in the system plan. .

The following example illustrates how the Integrated Virtualization Manager compares hardware resources in the system plan during the validation process to determine whether the system plan is valid for a managed system:

- A system plan specifies a server with two processors, 8 GB of memory, and a specific placement of physical I/O adapters within the system unit. If a server contains two processors, 16 GB of memory, a matching placement of physical I/O adapters within the system unit, and an expansion unit with additional physical I/O adapter, then the system passes validation.
- However, a server that contains 4 GB of memory causes the system plan to fail validation.
- A system plan also fails validation if the system plan specifies one type of physical I/O adapter in a slot but the actual system unit has a different type of physical I/O adapter in that slot. However, if the system plan specifies an empty slot, validation allows any type of physical I/O adapter to be in that slot on the actual system.

The Integrated Virtualization Manager also validates the disk drives that are attached to physical I/O adapters against the disk drives specified in the system plan.

When using a version of the Integrated Virtualization Manager prior to version 1.5.2, you can deploy system plans only to new systems, or to systems that do not already have the Integrated Virtualization Manager management partition configured. You can install the Integrated Virtualization Manager, but you cannot configure it.

If any step in the hardware validation process fails for the system plan, validation of the entire system plan fails.

The conditions that must be met for successful partition validation vary based on the version of the Integrated Virtualization Manager that you are using to deploy the system plan.

## Partition validation for Integrated Virtualization Manager prior to version 1.5.2

When using a version of Integrated Virtualization Manager prior to version 1.5.2 to deploy a system plan, the system plan and the target system must meet the following partition validation conditions:

1. The target system must be a new system or a system that does not already have the Integrated Virtualization Manager management partition configured. System plan deployment requires that you have Integrated Virtualization Manager installed, but not configured.
2. If the managed system that is the target of the system plan contains any of the following items, validation of the system plan will fail:
   - Client logical partitions
   - Virtual Ethernet adapters
   - Shared Ethernet adapters
   - EtherChannel adapters
   - Storage pools
   - Backing devices
3. If the managed system that is the target of the system plan does not have the following resources available as required by the system plan, the validation of the system plan will fail:
   - Processors
   - Memory
   - I/O adapters
   - Storage

If any step in the partition validation process fails for the system plan, validation of the entire system plan fails.

### Partition validation for Integrated Virtualization Manager version 1.5.2

When using Integrated Virtualization Manager version 1.5.2 to deploy a system plan, you can deploy a system plan to a system that is not new or that is not in the manufacturer default configuration.

When using Integrated Virtualization Manager version 1.5.2 to deploy a system plan, partition validation is more flexible. The partition validation criteria for Integrated Virtualization Manager version 1.5.2 has the following changes from previous versions:

- The Integrated Virtualization Manager management partition or client logical partitions can exist on the system that is the target of the system plan. However, any client logical partitions and hardware configured on the target system must be identical to those same items in the system plan.
- The target system can have Virtual I/O Server items, such as virtual disks or virtual Ethernet adapters, already configured. The items in the system plan and the items configured on the managed system need not match exactly. Based on certain criteria, the system plan validation process matches the Virtual I/O Server items in the plan and the items on the target system. If an item does not match, the item in the plan is a candidate for deployment. However, if the item in the system plan uses any adapter or device that some other configured item on the system is already using, the plan item is marked as not deployable. At the end of the validation process, the wizard presents a list of items in the system plan that are deployable and a list of items that are not deployable. When you continue the deployment of the system plan, the wizard does not attempt to deploy any items that are marked as not deployable.

  The following examples illustrate how the Integrated Virtualization Manager compares Virtual I/O Server items and other logical partition deployment items in the system plan during the validation process to determine whether these items can be deployed:

  - The system plan contains a shared Ethernet adapter for a logical partition, and the target system has no shared Ethernet adapters configured. In this case, the shared Ethernet adapter in the system plan is deployed to the target system.
  - The system plan contains two EtherChannel adapters, and an EtherChannel exists on the target system. The existing EtherChannel exactly matches one of the two that are in the system plan, including all physical adapters. In this case, this EtherChannel passes validation successfully, but is not deployed. The other EtherChannel in the system plan, which does not require any of the adapters that the existing EtherChannel or shared Ethernet adapter uses, is deployed.
  - The system plan contains two storage pools, and a storage pool exists on the target system. The existing storage pool exactly matches one of the two in the system plan, including all location codes for all disks and specified disk sizes. In this case, this storage pool passes validation successfully, but is not deployed. The other storage pool in the system plan, which has no attribute conflicts with the existing storage pool, is deployed.
  - The system plan contains an EtherChannel adapter, and a shared Ethernet adapter exists on the target system. The existing shared Ethernet adapter uses a physical adapter that is specified for use by the EtherChannel adapter in the system plan. In this case, the EtherChannel adapter fails validation and is cannot be deployed.
- 
- You no longer must deploy a system plan in its entirety, but can instead partially deploy a system plan on the target system by selecting which logical partitions in the plan to deploy. You can run the Deploy System Plan Wizard again at another time to deploy the remainder of the logical partitions in the system plan.

If any step in the partition validation process fails for the system plan, validation of the entire system plan fails.

## Installing the Integrated Virtualization Manager

Install the Virtual I/O Server management partition on an servers server or an BladeCenter blade server. Then, connect to the Integrated Virtualization Manager Web-based interface.

## Installing the Virtual I/O Server and enabling the Integrated Virtualization Manager on systems servers

When you install the Virtual I/O Server in an environment where no Hardware Management Console (HMC) is present, the Virtual I/O Server automatically creates a management partition whose interface is the Integrated Virtualization Manager.

Before you start, ensure that you have completed the following tasks:

1. Verify that you have cabled the server. Specifically, ensure that you have connected a serial cable from a PC or ASCII terminal to a system port on the server.
2. Verify that the system has a 512–byte disk drive available for the installation. The Virtual I/O Server only recognizes 512–byte disk drives.
3. Verify that you have access to the Advanced System Management Interface (ASMI) using the Web interface.
4. Verify that you have the Administrator or Authorized server provider authority level in ASMI.
5. Using the Web-based ASMI, change the following settings as appropriate for the type of partition on which you are installing the Integrated Virtualization Manager:

   For an AIX or Linux partition, complete the following steps to change the partition boot mode:

   a. In the navigation area, expand **Power/Restart Control**.

   b. Click **Power On/Off System**.

   c. Select **Boot to SMS menu** in the **AIX/Linux partition mode** boot field.

   d. Click **Save settings and power on**.
6. Open a terminal session on the PC, using an application such as HyperTerminal, and wait for the SMS menu to appear. Be sure the line speed is set to 19,200 bits per second to communicate with the system unit.
7. Using the Web-based ASMI, change the partition boot mode back so that the server loads the operating environment during startup:

   a. Expand **Power/Restart Control**.

   b. Click **Power On/Off System**.

   c. Select **Continue to operating system** in the **AIX/Linux partition mode** boot field.

   d. Click **Save settings**.

To install the Virtual I/O Server and enable the Integrated Virtualization Manager, complete the following steps:

1. Insert the Virtual I/O Server CD or DVD into the optical drive.
2. In SMS, select the CD or DVD as the boot device:

   a. Select **Select Boot Options**, and then press Enter.

   b. Select **Select Install/Boot Device**, and then press Enter.

   c. Select **CD/DVD**, and then press Enter.

   d. Select the media type that corresponds to the optical device, and then press Enter.

   e. Select the device number that corresponds to the optical device, and then press Enter.

   f. Select **Normal Boot**, and confirm that you want to exit SMS.
3. Install the Virtual I/O Server:

   a. Select the console, and then press Enter.

   b. Select a language for the BOS menus, and then press Enter.

   c. Select **Start Install Now with Default Settings**.

   d. Select **Continue with Install**. The managed system restarts after the installation is complete, and the login prompt is displayed on the ASCII terminal.

After you install the Integrated Virtualization Manager, finish the installation by accepting the license agreement, checking for updates, configuring the TCP/IP connection. For instructions, see "Finishing the Integrated Virtualization Manager installation" on page 171.

**Related information**

Installing the Virtual I/O Server using NIM

## Installing the Virtual I/O Server and enabling the Integrated Virtualization Manager on an BladeCenter blade server with POWER Architecture technology

When you install the Virtual I/O Server on an BladeCenter blade server with POWER Architecture technology, the Virtual I/O Server automatically creates a management partition whose interface is the Integrated Virtualization Manager.

Before you start, ensure that you have completed the following tasks:

1. Start a Telnet or SSH session to the BladeCenter management module. For instructions, see Telnet connection or Serial connection in the Information Center.
2. Start a Serial over LAN (SOL) session. For instructions, see Starting an SOL session in the Information Center.
3. Start the System Management Services (SMS) utility. For instructions, see Starting the SMS utility in the Information Center.

To install the Virtual I/O Server and enable the Integrated Virtualization Manager, complete the following steps:

1. Insert the Virtual I/O Server CD or DVD into the optical drive.
2. Assign the media tray to the blade server on which you plan to install the Virtual I/O Server:
   a. From the management module Web interface, select **Blade Tasks** → **Remote Control**.
   b. Select **Start Remote Control**.
   c. In the Change media tray owner field, select the blade server on which you plan to install the Virtual I/O Server.

   Alternatively, you can assign the media try to the blade server by using the control panel.
3. In SMS, select the CD or DVD as the boot device:
   a. Select **Select Boot Options**, and then press Enter.
   b. Select **Select Install/Boot Device**, and then press Enter.
   c. Select **List all Devices**, and then press Enter.
   d. Select the device number that corresponds to the optical device, and then press Enter.
   e. Select **Normal Boot Mode**, and then press Enter.
   f. Exit the SMS menu by pressing the x key, and confirm that you want to exit SMS.
4. Install the Virtual I/O Server:
   a. Select the console, and then press Enter.
   b. Select a language for the BOS menus, and then press Enter.
   c. Select **Change/Show Installation Settings and Install**, and then press Enter.
   d. Select **1** to verify that Disk Where You Want to Install field is set appropriately. Verify the actual location code (for example, 01-08-00-1,0) of the target hard disk. The logical name for the hard disks (for example, hdisk0) that is displayed in this menu can be different from the logical name for the same hard disk that is listed within the Virtual I/O Server (for example, from the lspv command) that runs on the same machine. This can happen when you add disks after you install the Virtual I/O Server.

e. Return to the Installation and Maintenance menu and select **Start Install Now with Default Settings**.

f. Select **Continue with Install**. The managed system restarts after the installation is complete, and the login prompt is displayed on the ASCII terminal.

After you install the Integrated Virtualization Manager, finish the installation by accepting the license agreement, checking for updates, and configuring the TCP/IP connection. For instructions, see "Finishing the Integrated Virtualization Manager installation."

**Related information**

Installing the Virtual I/O Server using NIM

## Finishing the Integrated Virtualization Manager installation

After you install the Integrated Virtualization Manager, you need to accept the license agreement, check for updates, configure the TCP/IP connection.

This procedure assumes that the Integrated Virtualization Manager is installed. For instructions, see one of the following tasks:

- "Installing the Virtual I/O Server and enabling the Integrated Virtualization Manager on systems servers" on page 169
- "Installing the Virtual I/O Server and enabling the Integrated Virtualization Manager on an BladeCenter blade server with POWER Architecture technology" on page 170

To finish the installation, complete the following steps:

1. Log in to the management partition with the user ID **padmin**.
2. When prompted, change the login password to a secure password that adheres to your local password-security guidelines.
3. Accept the Virtual I/O Server license agreement. For instructions, see "Viewing and accepting the Virtual I/O Server license" on page 72.
4. Ensure that there is a network connection configured between the management partition and at least one of the physical Ethernet adapters on the managed system. This allows you to access the Integrated Virtualization Manager interface from a computer that is connected to the physical Ethernet adapter. You cannot use the HMC1 and HMC2 ports to connect to the management partition.
5. Configure the TCP/IP connection for the Integrated Virtualization Manager by using the mktcpip command. You must configure TCP/IP before you can perform any dynamic logical partitioning operations. Integrated Virtualization Manager version 1.5.2, and later, supports the use of IPv6 addresses. To view the PDF file of the *Virtual I/O Server and Integrated Virtualization Manager Command Reference* (SA76-0101), approximately 4 MB in size, see sa76-0101.pdf .
6. Connect to the Web interface or the command-line interface. For instructions, see one of the following tasks:
   - "Connecting to the Integrated Virtualization Manager Web-based interface" on page 172
   - "Connecting to the Virtual I/O Server command-line interface" on page 172
7. Check for updates to the Integrated Virtualization Manager. For instructions, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.

When you are finished, configure the management partition and client logical partitions. For instructions, see "Configuring the management partition and client logical partitions" on page 172.

## Connecting to the Integrated Virtualization Manager Web-based interface

Learn how to connect to the Web-based system management interface for the Integrated Virtualization Manager.

You must know the IP address that is assigned to the Integrated Virtualization Manager.

To connect to the Web-based interface for the Integrated Virtualization Manager, do the following:

1. Open a Web browser window, and connect using the HTTP or HTTPS protocol to the IP address that was assigned to the Integrated Virtualization Manager during the installation process. For example, enter `https://123.456.7.890` in your Web browser, where *123.456.7.890* is the IP address assigned to the Integrated Virtualization Manager. The Welcome window is displayed.
2. Enter the default user ID of **padmin**, and enter the password that you defined during the installation process. The Integrated Virtualization Manager interface is displayed.

For information about the Web-based interface navigation, see the online help for the Integrated Virtualization Manager.

## Connecting to the Virtual I/O Server command-line interface

Learn how to connect to the Virtual I/O Server command-line interface, which allows you to use commands for the Integrated Virtualization Manager.

Connect to the Virtual I/O Server command-line interface using one of the following methods:

**Open a virtual terminal session to the management partition**
For instructions, see "Opening a virtual terminal session for a logical partition" on page 188.

**Telnet** You can use Telnet to connect to the command-line interface. Telnet does not provide a secure connection to the Virtual I/O Server. Therefore, use Telnet only if the Ethernet adapter that you have configured to access the management partition is physically isolated from networks that are not secure.

**OpenSSL or Portable OpenSSH**
You can use OpenSSL or Portable SSH to securely connect to the Virtual I/O Server from a remote location. For instructions, see "Connecting to the Virtual I/O Server using OpenSSH" on page 119.

# Configuring the management partition and client logical partitions

Find instructions for configuring the system by deploying a system plan or manually configuring the management partition and client logical partitions.

These instructions apply to configuring a system that is managed by the Integrated Virtualization Manager. If you plan to install the Virtual I/O Server on a system that is managed by a Hardware Management Console (HMC), then you need the instructions for installing and configuring the Virtual I/O Server on a system managed by an HMC. See "Installing the Virtual I/O Server and client logical partitions" on page 61.

The configuration procedures vary depending on whether you plan to deploy a system plan to configure the management partition and client logical partitions. When you deploy a system plan, the Integrated Virtualization Manager automatically configures virtual resources and creates the logical partitions and partition profiles for the client logical partitions based on the configuration specifications in the system plan.

To configure the management partition and client logical partitions by deploying a system plan, the Integrated Virtualization Manager must be at version 1.4.

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

You can enter the activation code for PowerVM Editions (or Advanced POWER Virtualization) using the Integrated Virtualization Manager.

The code level for the Integrated Virtualization Manager must be at version 1.5, or later, to perform the following procedure. For instructions about how to view and update the current code level, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.

Whether you need to enter an activation code depends on your edition of the PowerVM Editions feature and the hardware on which you plan to enable it. The following table summarizes the requirements.

*Table 48. Activation code requirements*

| | **systems** | **BladeCenter blade servers** |
|---|---|---|
| **PowerVM Express Edition** | No activation code is required. | The Express Edition is not available on blade servers. |
| **PowerVM Standard Edition** | The PowerVM Editions activation code is required. | No activation code is required. The Standard Edition is included with the blade server. |
| **PowerVM Enterprise Edition** | The PowerVM Editions activation code is required.<br>**Note:** If you already have the Standard Edition enabled, you must enter a separate, additional activation code for the Enterprise Edition. | The PowerVM Editions activation code is required. |

For detailed information about the PowerVM Editions editions, see "PowerVM Editions" on page 2.

Before you start, verify that you have access to the Integrated Virtualization Manager. For instructions, see "Installing the Virtual I/O Server and enabling the Integrated Virtualization Manager on systems servers" on page 169.

To enter the activation code in the Integrated Virtualization Manager, complete the following tasks:

1. From the **IVM Management** menu, click **Enter PowerVM Editions Key**. The Enter PowerVM Editions Key window is displayed.
2. Enter your activation code for PowerVM Editions and click **Apply**.

You can now create more than two client logical partitions that use virtual I/O or shared processors.

## Deploying a system plan by using the Integrated Virtualization Manager

When you deploy a system plan, the Integrated Virtualization Manager creates logical partitions on the managed system according to the specifications in the system plan.

**Requirements for deploying a system plan**

When you use a version of the Integrated Virtualization Manager prior to version 1.5.2.0, ensure that the system is in the manufacturing default configuration. More specifically, the system must meet the following requirements:

- Client logical partitions are not configured on the managed system.
- Virtual Ethernet adapters are not configured on the managed system.

- Storage pools are not configured on the managed system.
- Backing devices are not configured on the managed system.
- All of the I/O resources are assigned to the Integrated Virtualization Manager management partition. In the manufacturing default configuration, all the I/O resources are assigned to the Integrated Virtualization Manager. When you add I/O resources to the system, they are assigned to the Integrated Virtualization Manager by default.

When you deploy a system plan with Integrated Virtualization Manager version 1.5.2.0, you can deploy a system plan to a system that is not new or that is not in the manufacturer default configuration. You can deploy a system plan to a system that already has a configured Integrated Virtualization Manager management partition or that has configured client logical partitions. Also, the target system can have any of the previously listed Virtual I/O Server items, such as virtual disks or virtual Ethernet adapters, already configured. However, if the system plan that you intend to deploy contains information about any items that are already configured on the system, the configured items on the target system must exactly match those same items in the system plan. If they do not match exactly, then the system plan either cannot pass validation or the item in the system plan cannot be deployed.

In addition, you no longer must deploy a system plan in its entirety, but can instead partially deploy a system plan on the target system by selecting which logical partitions in the plan to deploy. You then can run the Deploy System Plan Wizard again to deploy the remainder of the logical partitions in the system plan at another time.

In addition to meeting the above requirements based on your version of the Integrated Virtualization Manager, you must meet the following prerequisites:

- The system-plan file exists on the Integrated Virtualization Manager. If the system-plan file does not exist on the Integrated Virtualization Manager, you must import the system-plan file into the Integrated Virtualization Manager.
- The physical hardware is connected and is reporting to the server. If you are deploying a system plan that you created by using the Integrated Virtualization Manager, verify that the hardware and cabling on the target system is identical to that on the source system.
- The physical hardware on the managed system must match exactly to any of the same hardware in the system plan.
- The Integrated Virtualization Manager is not performing any other operations on the managed system.
- You are the prime administrator (padmin). For more information about user roles, refer to User roles.

**Deploying a system plan**

To deploy a system plan on a managed system by using the Integrated Virtualization Manager, complete the following steps:

1. In the navigation area of the Integrated Virtualization Manager, select **Manage System Plans**. The Manage System Plans page opens.
2. In the System Plans table, select the system plan that you want to deploy.
3. Select **More Tasks** → **Deploy** from the toolbar at the top of the System Plans table to start the Deploy System Plan Wizard. The System Deployment: Deployment Overview page of the wizard opens.
4. If prompted, choose the managed system to which you want to deploy the system plan and click **Next**. The prompt only occurs if the system plan file contains more than one system. If the system plan does not match the hardware on the managed system to which you want to deploy the plan, the wizard displays a window that informs you of this. Click **OK** to continue or **Cancel** to select a different system plan.
5. Wait for the wizard to validate the managed system and its hardware against the system plan. The validation process can take several minutes.
6. If the validation process completes successfully, click **Next**. If the validation process does not complete successfully, correct the issues indicated by the error messages, click **Cancel** to exit the wizard, and

restart this procedure from the beginning. To help you to correct validation issues, you might want to create a system plan that is based on the current configuration of the managed system. Such a system plan can help you to compare the system plan that you want to deploy with the current configuration of the managed system. You can do this by using the Create System Plan task in the Integrated Virtualization Manager, or you can run the following command on the system:

```
mksysplan -f name_of_new_system_plan.sysplan
```

7. Review the Deployable Plan Items page, select the logical partitions in the system plan that you want to deploy, and click **Next**. This page indicates the deployable status of the logical partitions that the system plan contains. If a logical partition has a status of partially deployed, the logical partition is selected for deployment automatically, and you cannot change the selection. If a logical partition is deselected for deployment, the wizard does not deploy any dependent entities for that partition, such as backing devices on the Virtual I/O Server.
8. Review the list of deployment plan items on the Deployment page, and click **Finish** to begin deploying the system plan. The Integrated Virtualization Manager creates the specified logical partitions and deploys the specified entities as listed. The deployment process can take several minutes depending on the number of logical partitions and entities to be deployed.

After you finish the deployment of the system plan, complete the following tasks:

- Locate the physical disk I/O adapters that belong to each logical partition and verify that the disk drives that are attached to these physical I/O adapters will support your desired configuration for each logical partition.
- Install operating systems and software on the logical partitions.

**Related concepts**

"User roles" on page 202
Learn about the user roles for the Integrated Virtualization Manager.

**Related tasks**

"Importing a system plan into the Integrated Virtualization Manager" on page 195
You can import a system-plan file into the Integrated Virtualization Manager management partition. You can then deploy the system plan to the system that the Integrated Virtualization Manager manages.

## Manually configuring the management partition and client logical partitions

You can configure virtual resources on the management partition and create the client logical partitions and partition profiles.

Before you begin, complete the following tasks:

- Determine the system resource requirements for the Virtual I/O Server management partition. The system resource requirements for the management partition can depend on many factors. These factors can include the server model, the number of logical partitions you create on the managed system, and the number of virtual devices used by those logical partitions.

  When you install the Virtual I/O Server, it automatically creates a logical partition for itself on the server. (This logical partition is called the *management partition*.) The Virtual I/O Server automatically assigns a fraction of the memory and processors on the server to the management partition. You can change the default amount of memory and processor resources that are assigned to the management partition.

  For instructions, see "Planning for the Virtual I/O Server" on page 45.
- Develop a plan for the storage needs of each logical partition that you are to create on your managed system. Calculate how much storage space each logical partition requires for its operating system, applications, and data. For more information about the storage requirements for each operating system, consult the operating system documentation.

## Changing memory and processor resources on the management partition

Use the Integrated Virtualization Manager to change the memory and processor resources on the management partition.

Before you start, complete the following tasks:

1. Install the Integrated Virtualization Manager. For instructions, see "Installing the Integrated Virtualization Manager" on page 168.
2. Ensure that your user role is not View Only.

If you want to change the memory and processor resources on the management partition, do the following:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the management partition (**partition ID 1**).
3. From the **Tasks** menu, click **Properties**. The Partition Properties panel is displayed.
4. Click the **Memory** tab to display the memory settings.
5. Change the minimum, assigned, and maximum pending amounts of memory to the amount of memory that you want the management partition to use. If you are using a workload-management application, then you can set the minimum and maximum amounts for the management partition. The assigned amount is the amount of memory that the management partition has initially assigned to it. If you do not anticipate dynamically increasing the memory beyond a certain point, setting the appropriate maximum value will save on reserved firmware memory.
6. Click the **Processing** tab to display the processing settings. Keep the default settings unless you are using a workload-management application.
7. Click **OK** to apply the changes. It might take a few minutes for the managed system to apply the changes. If you changed a minimum or maximum value, restart the system for the changes to take effect.

## Setting the maximum number of logical partitions

Set the maximum number of logical partitions that you want to allow on the managed system using the Integrated Virtualization Manager.

Before you start, ensure that your user role is not View Only.

You can set the maximum number of logical partitions that you want to allow on this managed system. The managed system reserves a small amount of system memory to accommodate the maximum number of logical partitions that you specify.

To set the maximum number of logical partitions, do the following:

1. From the **Partition Management** menu, click **View/Modify System Properties**. The View/Modify System Properties panel is displayed.
2. In the **Configured maximum** field, verify whether this is the maximum number of logical partitions that you want to allow on this managed system. If it is not, then do the following:
   a. Specify the maximum number of logical partitions in the **Maximum after restart** field, and click **OK**.
   b. Open a virtual terminal session to the management partition. For instructions, see "Opening a virtual terminal session for a logical partition" on page 188.
   c. Restart the system. For instructions, see "Shutting down logical partitions" on page 188. It might take a few minutes for the managed system to restart. Be sure to complete all of the setup steps before restarting the system. Otherwise, you might need to restart the system more than once.

## Mirroring the Integrated Virtualization Manager management partition

To prevent potential downtime and data loss, add a second disk to the rootvg storage pool and mirror the two disks.

When you install the Virtual I/O Server, Virtual I/O Server automatically creates a storage pool called rootvg and assigns one physical volume to rootvg. The Virtual I/O Server software (including the Integrated Virtualization Manager) and any data that the Virtual I/O Server software uses initially is stored on the management partition (partition ID 1) on that physical volume. If that disk would fail, you would be unable to manage your client partitions and would suffer downtime and the loss of data. To prevent this kind of interruption to your business, you need to add a second disk to the rootvg storage pool and mirror the two disks.

Before you start, ensure you meet the following requirements:

1. The Integrated Virtualization Manager is at version 1.5 or later. To update the Integrated Virtualization Manager, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.
2. You are the prime administrator (padmin).

To mirror the management partition, complete the following steps:

1. Add a new physical volume to the rootvg storage pool. For instructions, see "Modifying storage pools using the Integrated Virtualization Manager" on page 191.
2. To mirror the new volume to ensure that the it has all of the software and data that the original volume has, complete the following steps:
   a. Open a virtual terminal window to the management partition. For instructions, see "Opening a virtual terminal session for a logical partition" on page 188.
   b. Sign on to Virtual I/O Server using the padmin user ID and password.
   c. At the command prompt, run the mirriorios command as follows:

      ```
      mirrorios Physicalvolume
      ```

      where *Physicalvolume* is the name of the volume that you just added to rootvg.

      **Restriction:** The mirrorios command mirrors only the rootvg storage pool. It does not mirror other volume groups or any virtual disks that are created on rootvg after it is initially mirrored.

## Configuring storage on the managed system using the Integrated Virtualization Manager

You can create a storage pool in addition to the default storage pool, add additional physical volumes to the default storage pool, and create virtual disks using the Integrated Virtualization Manager.

You can assign storage to logical partitions in the following ways:

- You can assign physical volumes directly to the logical partition. (A *physical volume* is an individual logical unit that is identified by a logical unit number (LUN). A physical volume can be a hard disk or a logical device on a storage area network (SAN).)
- You can add physical volumes or files to a storage pool, create virtual disks from the storage capacity of the storage pool, and assign the virtual disks to logical partitions. Virtual disks allow you to specify more precisely the amount of storage that you assign to logical partitions. You can assign storage to logical partitions without regard to the actual capacities of the physical volumes or files that make up the storage pool.

Consider creating a storage pool in addition to the default rootvg storage pool for regular data storage, and then assign the new storage pool as the default. You can then add more physical volumes to a storage pool, create virtual disks from a storage pool, and assign these virtual disks to other logical partitions.

If you plan to assign physical volumes directly to logical partitions, you do not need to do anything with the physical volumes. You can assign the physical volumes to the logical partitions when you create the logical partitions.

To configure storage on the managed system, complete the following steps:

1. Create a second storage pool for regular data storage. For instructions, see “Creating storage pools.”
2. Add additional physical volumes to the default storage pool. For instructions, see “Modifying storage pools using the Integrated Virtualization Manager” on page 191.
3. Create virtual disks from the default storage pool. For instructions, see “Creating virtual disks.”

**Creating storage pools:**

You can create a logical volume based or file based storage pool on your managed system using the Integrated Virtualization Manager.

To create a logical volume based storage pool, you must assign at least one physical volume to the storage pool. When you assign physical volumes to a storage pool, the managed system erases the information on the physical volumes, divides the physical volumes into physical partitions, and adds the capacity of the physical partitions to the storage pool. Do not add a physical volume to the storage pool if the physical volume contains data that you want to preserve.

To create file based storage pools, the Integrated Virtualization Manager must be at version 1.5 or later. To update the Integrated Virtualization Manager, see “Viewing and updating the code level of the Integrated Virtualization Manager management partition” on page 199.

To create a storage pool, do the following:

1. From the **Virtual Storage Management** menu, click **View/Modify Virtual Storage**.
2. Click the **Storage Pools** tab.
3. Click ***Create Storage Pool...**. The Create Storage Pool window displays.
4. Enter a name for the storage pool and select the storage pool type.
5. Enter or select the information required for the logical volume based or file based storage pool and click **OK**.

**Creating virtual disks:**

Use the Integrated Virtualization Manager to create a virtual disk on your managed system. Virtual disks are also known as *logical volumes*.

To create a virtual disk, do the following:

1. From the **Virtual Storage Management** menu, click **View/Modify Virtual Storage**.
2. Click ***Create Virtual Disk...**. The Create Virtual Disk panel is displayed.
3. Enter a virtual disk name, select a storage pool, and enter a size for the virtual disk, and then click **OK**. The virtual disk is created, and then the View/Modify Virtual Storage panel is displayed.
4. Repeat this procedure for each virtual disk that you want to create.
5. To view or modify the properties of any virtual disks that you just created, see “Modifying virtual disks” on page 191.

These steps are equivalent to using the **mkbdsp** command in the command-line interface.

If there is not enough disk space for the virtual disk, increase the size of the default storage pool. For instructions, see “Modifying storage pools using the Integrated Virtualization Manager” on page 191

## Configuring Ethernet on the managed system by using the Integrated Virtualization Manager

You can create virtual Ethernet bridges, configure a Host Ethernet Adapter (or Integrated Virtual Ethernet), and assign physical Ethernet adapters to client logical partitions by using the Integrated Virtualization Manager.

You can configure the following types of Ethernet on your managed system:

- You can create virtual Ethernet bridges on your managed system. Virtual Ethernet bridges, also known as Shared Ethernet Adapters, connect the virtual Ethernet networks on your managed system to physical Local Area Networks (LANs). For greater security, do not set up the physical Ethernet adapter that you use to connect to the management partition as a virtual Ethernet bridge. This allows you to isolate the management partition from all external networks. (The management partition manages the virtual Ethernet networks on your managed system but does not participate in any virtual Ethernet networks.)

  If you configure a single physical Ethernet adapter or link aggregation to connect to the management partition and to act as a virtual Ethernet bridge, consider using OpenSSL and Portable OpenSSH on the management partition. You can use OpenSSL and Portable OpenSSH to connect securely to the Virtual I/O Server from a remote location.

  You do not need to select a physical Ethernet adapter or link aggregation for a virtual Ethernet network. If no physical adapter or link aggregation is set for a virtual Ethernet network, then the logical partitions on the virtual Ethernet network can communicate with one another, but they cannot communicate directly with a physical network.
- You can configure a Host Ethernet Adapter port. A Host Ethernet Adapter is a unique Ethernet adapter that is built into the system. It provides the ability to partition a physical Ethernet port. A Host Ethernet Adapter can contain one or more physical ports and each physical port can be assigned to zero or more logical partitions.
- You can assign a physical Ethernet adapter to a client logical partition.

To configure virtual Ethernet on the managed system, complete the following steps:

1. Configure virtual Ethernet bridges. For instructions, see "Configuring virtual Ethernet bridges on the managed system using the Integrated Virtualization Manager."
2. Configure a Host Ethernet Adapter. For instructions, see "Assigning a Host Ethernet Adapter port to a logical partition" on page 180
3. Assign a physical Ethernet adapter to a client logical partition. For instructions, see "Dynamically managing physical adapters" on page 180.

### Configuring virtual Ethernet bridges on the managed system using the Integrated Virtualization Manager:

Use the Integrated Virtualization Manager to configure virtual Ethernet bridges on the managed system.

A physical Ethernet adapter or link aggregation that connects a virtual Ethernet network with a physical local area network (LAN) is called a *virtual Ethernet bridge*. Another name for a virtual Ethernet bridge is a *shared Ethernet adapter* because the logical partitions on the virtual Ethernet network share the physical Ethernet connection. Virtual Ethernet bridges connect the virtual Ethernet networks on your managed system to physical LANs.

For greater security, do not set up the physical Ethernet adapter or link aggregation that you use to connect to the management partition as a virtual Ethernet bridge. This situation allows you to isolate the management partition from all external networks. (The management partition manages the virtual Ethernet networks on your managed system, but it does not participate in any virtual Ethernet networks.)

If you configure a single physical Ethernet adapter or link aggregation to connect to the management partition and to act as a virtual Ethernet bridge, consider installing OpenSSL and Portable OpenSSH on the management partition. You can use OpenSSL and Portable OpenSSH to connect securely to the Virtual I/O Server from a remote location.

You do not need to select a physical Ethernet adapter or link aggregation for a virtual Ethernet network. If no physical adapter or link aggregation is set for a virtual Ethernet network, then the logical partitions on the virtual Ethernet network can communicate with one another, but they cannot communicate directly with a physical network.

Your role must not be View Only or Service Representative (SR) to perform this procedure.

To configure virtual Ethernet bridges, do the following:

1. From the **Virtual Ethernet Management** menu, click **View/Modify Virtual Ethernet**. The View/Modify Virtual Ethernet panel is displayed.
2. Click the **Virtual Ethernet Bridge** tab.
3. Set each **Physical Adapter** field to the physical adapter that you want to use as the virtual Ethernet bridge for each virtual Ethernet network. (The HMC1 and HMC2 ports do not display in the **Physical Adapter** field and cannot be used as virtual Ethernet bridges.)
4. Click **Apply** to apply the changes.

**Assigning a Host Ethernet Adapter port to a logical partition:**

Use the Integrated Virtualization Manager to assign a Host Ethernet Adapter (or Integrated Virtual Ethernet) port to a logical partition so that the logical partition can directly access the external network.

Before you start, ensure that the Integrated Virtualization Manager is at version 1.5 or later. To update the Integrated Virtualization Manager, see “Viewing and updating the code level of the Integrated Virtualization Manager management partition” on page 199.

A Host Ethernet Adapter is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. Host Ethernet Adapters offer high throughput, low latency, and virtualization support for Ethernet connections.

Unlike most other types of I/O devices, you can never assign the Host Ethernet Adapter itself to a logical partition. Instead, multiple logical partitions can connect directly to the Host Ethernet Adapter and use the Host Ethernet Adapter resources. This allows these logical partitions to access external networks through the Host Ethernet Adapter without having to go through an Ethernet bridge on another logical partition.

To assign a Host Ethernet Adapter port to a logical partition, complete the following steps:

1. From the **I/O Adapter Management** menu, click **View/Modify Host Ethernet Adapters**.
2. Select a port with at least one available connection and click **Properties**.
3. Select the **Connected Partitions** tab.
4. Select the logical partition that you want to assign to the Host Ethernet Adapter port and click **OK**. If you want to remove a partition assignment, deselect the logical partition and click **OK**.

You also can use the Performance area of the **General** tab to adjust the settings for the selected Host Ethernet Adapter port. You can view and modify the speed, maximum transmission unit, and other settings for the selected port.

## Dynamically managing physical adapters

You can add and remove physical adapters to and from a running logical partition.

You can change the physical adapter settings for a logical partition at any time if the partition is capable of dynamic I/O adapter changes.

When making dynamic I/O adapter changes, keep the following items in mind:

- You might lose data if you remove a physical adapter from a running logical partition.
- You cannot assign a physical adapter to another partition if it is being used by the operating system of the partition to which it is currently assigned. If you attempt to reassign the adapter, an error message is displayed. You must unconfigure the device by using the tools of the appropriate operating system before you can change the adapter's partition assignment.

Before you start, ensure that the Integrated Virtualization Manager is at version 1.5 or later. To update the Integrated Virtualization Manager, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.

To dynamically add or remove physical adapters to or from a running logical partition, follow these steps:

1. If no client logical partitions exist, go to step 4.
2. Select the logical partition to which you want to assign a physical adapter and click **Properties**.
3. Verify that **Yes** is displayed for **I/O adapter DLPAR Capable**. You might need to click **Retrieve Capabilities** to verify this value. If **No** is displayed for **Processing DLPAR Capable**, then you cannot dynamically add or remove physical adapters to or from the logical partition.
4. From the **I/O Adapter Management** menu, click **View/Modify Physical Adapters**.
5. Select the adapter whose partition assignment you want to change and click **Modify Partition Assignment**.
6. Select the logical partition to which you want to assign the physical adapter and click **OK**. If you want to make this adapter available to any client logical partition, including those not yet created, select **None** as the **New partition**.

## Creating client logical partitions using the Integrated Virtualization Manager

You can create client logical partitions on the managed system by deploying a system plan, using the Create Partitions wizard, or creating partitions based on existing partitions.

To create client logical partitions on your managed system, complete one of the following steps:

- Create client logical partitions by deploying a system plan. For instructions, see "Deploying a system plan by using the Integrated Virtualization Manager" on page 173.
- Create client logical partitions using the Create Partitions wizard. For instructions, see "Creating client logical partitions using the Create Partitions wizard."
- Create client logical partitions based on existing client logical partitions. For instructions, see "Creating a partition based on an existing partition" on page 182.

When you are finished, you can activate the client logical partitions and install their operating systems. For instructions, see the following information:

- "Activating logical partitions" on page 183
- AIX Installation and Migration
- Installing Linux

### Creating client logical partitions using the Create Partitions wizard:

You can use the Create Partitions wizard on the Integrated Virtualization Manager to create a new client logical partition on your managed system.

Use any role other than View Only to perform this task. However, do not use the Service Representative (SR) user role for this task because it cannot configure the storage in the Create Partition wizard.

To create a logical partition on your managed system, do the following:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Click ***Create Partition...** The Create Partition wizard is displayed.
   a. Follow the instructions on each step of the wizard, and then click **Next** when you have completed each step.
   b. When the Summary step is displayed, confirm that the information displayed in this step is correct, and then click **Finish**.
3. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed; the new partition is listed.

**Creating a partition based on an existing partition:**

Use the Integrated Virtualization Manager to create a new logical partition that is based on an existing partition on your managed system.

Use any role other than View Only to perform this task.

Use this task to create a new logical partition with the same properties as the selected existing partition with the exception of ID, name, physical volumes, and optical devices.

To create a logical partition based on an existing partition, do the following steps:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition that you want to use as a basis for the new partition.
3. From the **Tasks** menu, click **Create based on**. The Create Based On panel is displayed.
4. Enter the name of the new partition and determine whether you want to create virtual disks with the new partition.
5. Click **OK**. The View/Modify Partitions panel is displayed; the new partition is listed.

## Managing the system with the Integrated Virtualization Manager

You can manage all aspects of the system with the Integrated Virtualization Manager including modifying processor, memory, networking, and storage resources across the logical partitions on the system.

## Viewing and modifying system properties

Use the Integrated Virtualization Manager to view and modify the properties that apply to your managed system in general.

Use any role other than View Only to perform this task. The View Only role can view the properties, but it cannot modify them.

To view and modify your system properties, do the following:

1. From the **Partition Management** menu, click **View/Modify System Properties**. The View/Modify System Properties panel is displayed.
2. Depending on which properties you want to view and modify, click one of the following tabs:
   - **General** to view and modify the information that identifies this managed system, system status, and maximum number of logical partitions
   - **Memory** to view and modify the memory usage information for your managed system in general
   - **Processing** to view and modify processor usage information for your managed system in general

For more information about specific system properties that you can view or modify, see the online help ( ? ).

## Managing partitions using the Integrated Virtualization Manager

Use the logical partition management tasks to create and manage the logical partitions on your managed system with the Integrated Virtualization Manager.

### Activating logical partitions

Use the Integrated Virtualization Manager to activate logical partitions on the managed system.

Use any role other than View Only to perform this task.

You can activate logical partitions manually after you power on the managed system, or you can reactivate a logical partition after you have shut down the logical partition manually.

To activate a logical partition, do the following:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition you want to activate. You can select more than one partition at a time.
3. Click **Activate**. The Activate Partitions panel is displayed. Verify the partition ID, partition name, and the current state of the logical partition.
4. Click **OK** to activate the partition. The View/Modify Partitions panel is displayed, and the partition is activated.

Each logical partition is activated with the boot mode and keylock position that are selected on the Partition Properties panel for the logical partition.

For more information about activating logical partitions, see the online help ( ? ).

### Adding a client logical partition to the partition workload group

If you want to manage logical partition resources using a workload management tool, then you need to add the client logical partition to the partition workload group.

A *partition workload group* identifies a set of logical partitions that are located on the same physical system. Workload management tools use partition workload groups to identify which logical partitions they can manage. For example, Enterprise Workload Manager™ (EWLM) can dynamically and automatically redistribute processing capacity within a partition workload group to satisfy workload performance goals. EWLM adjusts processing capacity based on calculations that compare the actual performance of work processed by the partition workload group to the business goals defined for the work.

Workload management tools use dynamic logical partitioning (DLPAR) to make resource adjustments based on performance goals. For example, the partition management function of EWLM adjusts processor resources based on workload performance goals. Thus, EWLM can adjust the processing capacity for AIX and Linux logical partitions.

**Limitations:**

- Do not add the management partition to the partition workload group. To manage logical partition resources, workload management tools often require that you install some type of management or agent software on the logical partitions. To avoid creating an unsupported environment, do not install additional software on the management partition.
- For AIX and Linux partitions, the DLPAR support of the operating system is not the same as the DLPAR capabilities that are in the partition properties for a logical partition. The DLPAR support of the operating system reflects what each operating system supports with regard to DLPAR functions.

AIX and Linux support DLPAR of processors, memory, and I/O. The DLPAR capabilities that are shown in the partition properties for a logical partition reflect a combination of the following:

  - A Resource Monitoring and Control (RMC) connection between the management partition and the client logical partition
  - The operating system's support of DLPAR

  For example, an AIX client logical partition does not have an RMC connection to the management partition, but AIX supports DLPAR of processors, memory, and I/O. In this situation, the DLPAR capabilities shown in the partition properties for the AIX logical partition indicate that the AIX logical partition is not capable of processor, memory, or I/O DLPAR. However, because AIX supports DLPAR of processors, memory, and I/O, a workload management tool can dynamically manage its resources. Workload management tools are not dependent on RMC connections to dynamically manage logical partition resources.

- If a logical partition is part of the partition workload group, you cannot dynamically manage its resources from the Integrated Virtualization Manager because the workload management tool is in control of dynamic resource management. Not all workload management tools dynamically manage processor, memory, and I/O resources. When you implement a workload management tool that manages only one resource type, you limit your ability to dynamically manage the other resource types. For example, EWLM dynamically manages processor resources, but not memory or I/O. AIX supports processor, memory, and I/O DLPAR. EWLM controls dynamic resource management of processor resources, memory, and I/O for the AIX logical partition, but EWLM does not dynamically manage memory or I/O. Because EWLM has control of dynamic resource management, you cannot dynamically manage memory or I/O for the AIX logical partition from the Integrated Virtualization Manager.

To add a logical partition to the partition workload group, complete the following steps:

1. From the Partition Management menu, click **View/Modify Partitions**. The View/Modify Partitions window displays.
2. Select the logical partition that you want to include in the partition workload group.
3. From the Tasks menu, select **Properties**. The **Partition Properties** window is displayed.
4. In the General tab, select **Partition workload group participant** and click **OK**.

**Related information**

Enabling partition management

## Deleting logical partitions

Use the Integrated Virtualization Manager to delete logical partitions from the managed system.

Use any role other than View Only to perform this task.

When you delete a logical partition, all memory, processor, and storage resources that belonged to the logical partition become available for assignment to other logical partitions.

To delete a logical partition, do the following:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition that you want to delete.
3. From the **Tasks** menu, click **Delete**. The Delete Partitions panel is displayed, which shows the partition ID and partition name that identify each logical partition within the managed system and the current state of each logical partition. There is also an option to delete associated virtual disks for the partition.
4. Click **OK** to delete the partition. The View/Modify Partitions panel is displayed, and the partition is deleted.

Each logical partition is activated with the boot mode and keylock position that are selected on the Partition Properties panel for the logical partition.

For more information about deleting logical partitions, see the online help (?).

## Dynamically managing memory

Learn how to add and remove memory to and from a running logical partition.

You can add or remove memory for a running logical partition if the partition is capable of dynamic memory changes.

To dynamically add or remove memory to or from a running logical partition, follow these steps:

1. Select the logical partition that you want to change and click **Properties**.
2. Verify that **Yes** is displayed for **Memory DLPAR Capable**. You might need to click **Retrieve Capabilities** to verify this value. If **No** is displayed for **Memory DLPAR Capable**, then you cannot dynamically add or remove memory to or from the logical partition.
3. Click the **Memory** tab.
4. Specify new values in the Pending column.
5. Click **OK**. The management partition synchronizes the current assigned value with the pending assigned value. Synchronization can take several seconds to complete. You can perform other tasks on the system while the management partition is synchronizing the current and pending values. For more information about pending values and monitoring synchronization, see the online help (?).

## Dynamically managing physical adapters

You can add and remove physical adapters to and from a running logical partition.

You can change the physical adapter settings for a logical partition at any time if the partition is capable of dynamic I/O adapter changes.

When making dynamic I/O adapter changes, keep the following items in mind:

- You might lose data if you remove a physical adapter from a running logical partition.
- You cannot assign a physical adapter to another partition if it is being used by the operating system of the partition to which it is currently assigned. If you attempt to reassign the adapter, an error message is displayed. You must unconfigure the device by using the tools of the appropriate operating system before you can change the adapter's partition assignment.

Before you start, ensure that the Integrated Virtualization Manager is at version 1.5 or later. To update the Integrated Virtualization Manager, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.

To dynamically add or remove physical adapters to or from a running logical partition, follow these steps:

1. If no client logical partitions exist, go to step 4 on page 181.
2. Select the logical partition to which you want to assign a physical adapter and click **Properties**.
3. Verify that **Yes** is displayed for **I/O adapter DLPAR Capable**. You might need to click **Retrieve Capabilities** to verify this value. If **No** is displayed for **Processing DLPAR Capable**, then you cannot dynamically add or remove physical adapters to or from the logical partition.
4. From the **I/O Adapter Management** menu, click **View/Modify Physical Adapters**.
5. Select the adapter whose partition assignment you want to change and click **Modify Partition Assignment**.

6. Select the logical partition to which you want to assign the physical adapter and click **OK**. If you want to make this adapter available to any client logical partition, including those not yet created, select **None** as the **New partition**.

## Dynamically managing processing power

Learn how to add and remove processing power to and from a running logical partition.

You can add or remove processing power for a running logical partition if the partition is capable of dynamic processing power changes.

To dynamically add or remove processing power to or from a running logical partition, follow these steps:

1. Select the logical partition that you want to change and click **Properties**.
2. Verify that **Yes** is displayed for **Processing DLPAR Capable**. You might need to click **Retrieve Capabilities** to verify this value. If **No** is displayed for **Processing DLPAR Capable**, then you cannot dynamically add or remove processing power to or from the logical partition.
3. Click the **Processing** tab.
4. Specify new values in the Pending columns for Processing Units, Virtual Processors, and Uncapped weight.
5. Click **OK**. The management partition synchronizes the current assigned value with the pending assigned value. Synchronization can take several seconds to complete. You can perform other tasks on the system while the management partition is synchronizing the current and pending values. For more information about pending values and monitoring synchronization, see the online help (?).

## Modifying partition properties

Use the Integrated Virtualization Manager to view and modify the properties of the logical partition.

Use any role other than View Only to perform this task. Users with the Service Representative (SR) user role cannot view or modify storage values.

If the logical partition is powered off, then you can use this procedure to change many of the logical partition properties. The changes take effect when you reactivate the logical partition.

To view and modify the properties of the logical partition, do the following:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition for which you want to view or modify the properties.
3. From the **Tasks** menu, click **Properties**. The Partition Properties panel is displayed.
4. Click **OK** to activate the partition. The View/Modify Partitions panel is displayed, and the partition is activated.
5. Depending on which properties you want to view and modify, click one of the following tabs:
   - **General** to view the logical partition identifiers and the operating state. For AIX and Linux partitions, you can view or change certain identifiers and startup information, including boot mode and keylock position. You can also view and change dynamic LPAR (DLPAR) information, such as the partition host name or IP address, partition communication state, and the DLPAR capabilities of the partition.
   - **Memory** to view or modify the memory management information for the logical partition you selected.
   - **Processing** to view or modify the processor management settings for the logical partition you selected. For example, you can view the processor compatibility mode and set your preference for idle processor sharing for dedicated partitions.

- **Ethernet** to view or modify the logical partition settings for Host Ethernet Adapters (or Integrated Virtual Ethernet), virtual Ethernet adapters, and physical Ethernet adapters. The Ethernet settings that you can modify vary based on the operating system for the selected partition.
- **Storage** to view or modify the logical partition storage settings.
- **Optical Devices** to view or modify the logical partition settings for physical optical devices and virtual optical devices.
- **Physical adapters** to view or modify the physical adapters assigned to each logical partition.

The **Storage** and **Optical Devices** tabs are displayed for all logical partitions except the management partition.

For more information about specific partition properties that you can view or modify, see the online help (?).

## Migrating a client logical partition to another managed system

You can migrate an inactive or running client logical partition to another system managed by another Integrated Virtualization Manager.

When both the source and destination systems are managed by the Integrated Virtualization Manager at version 1.5, or later, you can migrate partitions between JS/22 Express systems and between JS/12 Express systems.

When both the source and destination systems are managed by the Integrated Virtualization Manager at version 1.5, or later, you can migrate AIX or Linux partitions between systems as follows.

*Table 49. Source and destination servers for partition migration with the Integrated Virtualization Manager*

| Source server | Destination server |
|---|---|
| servers POWER6-processor based server | servers POWER6-processor based server |
| servers POWER6-processor based server | JS/22 Express system or JS/12 Express system |
| JS/22 Express system or JS/12 Express system | JS/22 Express system or JS/12 Express system |
| JS/22 Express system or JS/12 Express system | servers POWER6-processor based server |

Before you start, complete the following tasks:

1. Ensure that the PowerVM Enterprise Edition (or Advanced POWER Virtualization Enterprise Edition) is enabled on both the source and destination servers. For instructions, see “Entering the activation code for PowerVM Editions with the Integrated Virtualization Manager” on page 10.
2. Ensure that you have properly prepared the source and destination systems and the migrating partition for the migration. For instructions, see “Preparing for an Integrated Virtualization Manager migration” on page 260..
3. Retrieve the IP address or host name of the Integrated Virtualization Manager that manages the system to which you plan to migrate the partition.

You cannot migrate the management partition.

To migrate a client logical partition to another managed system, complete the following steps:

1. From the View/Modify Partitions panel, select the client logical partition that you want to migrate, and click **Migrate**.
2. Enter the requested information and click **Validate**.
3. If you receive validation errors, fix the errors and return to this panel.
4. when you have fixed all validation errors, click **Migrate**.

To view the progress of the migration, see “Viewing or modifying the status of a migrating partition” on page 189

## Opening a virtual terminal session for a logical partition

Connect to a logical partition using the virtual terminal on the Integrated Virtualization Manager.

Before you start, ensure that the Integrated Virtualization Manager is at version 1.5 or later. To update the Integrated Virtualization Manager, see “Viewing and updating the code level of the Integrated Virtualization Manager management partition” on page 199.

To open a virtual terminal session, complete the following steps:

1. From the **Partition Management** menu, click **View/Modify Partitions**.
2. Select the logical partition to which you want to connect and click **Open terminal window**. A virtual terminal window displays.

   **Note:** Because the applet has a digital signature, your browser might display a security warning and ask you to verify that you want to run the applet.
3. Enter your password for your login ID from the current Integrated Virtualization Manager. A terminal session starts for the partition.

## Shutting down logical partitions

Use the Integrated Virtualization Manager to shut down the selected logical partitions or the entire managed system.

Use any role other than View Only to perform this task.

The Integrated Virtualization Manager provides the following types of shutdown options for logical partitions:

- Operating System (recommended)
- Delayed
- Immediate

The recommended shutdown method is to use the client operating systems shutdown command. Using the immediate shutdown method should be used as a last resort as this causes an abnormal shutdown which might result in data loss.

If you choose the Delayed shutdown method, then be aware of the following considerations:

- Shutting down the logical partitions is equivalent to pressing and holding the white control-panel power button on a server that is not partitioned.
- Use this procedure only if you cannot successfully shut down the logical partitions through operating system commands. When you use this procedure to shut down the selected logical partitions, the logical partitions wait a predetermined amount of time to shut down. This allows the logical partitions time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it ends abnormally, and the next restart might take a long time.

If you plan to shut down the entire managed system, shut down each client logical partition and then shut down the Virtual I/O Server management partition.

To shut down a logical partition, do the following:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition that you want to shut down.
3. Click **Shutdown**. The Shutdown Partitions panel is displayed.

4. Select the shutdown type.
5. Optional: Select **Restart after shutdown completes** if you want the logical partition to start immediately after it shuts down.
6. Click **OK** to shut down the partition. The View/Modify Partitions panel is displayed, and the partition is shut down.

For more information about shutting down logical partitions, see the online help (?).

## Using the operator panel service functions

This topic describes how to shut down, restart, or initiate a system memory dump on logical partitions using operator panel service functions in the Integrated Virtualization Manager. These functions are also known as *control panel functions*.

Use any role other than View Only to perform this task.

You can use the operator panel service functions to shut down or restart a logical partition without shutting down the operating system of that logical partition first.

**Attention:** Use this procedure only if you cannot successfully shut down or restart the logical partition through operating system commands. These operator panel service functions cause the logical partition to shut down abnormally and can cause data loss. The programs running in those processes are not allowed to perform any cleanup. These functions can cause undesirable results if data has been partially updated.

To use the operator panel service functions, complete the following steps:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition on which you want to perform the function.
3. From the **Tasks** menu, click **Operator panel service functions**. The Operator Panel Service Functions panel is displayed.
4. Select the operator panel service function that you want to use for the selected logical partition, and then click **OK**. The View/Modify Partitions panel is displayed, and the logical partition is shut down or restarted.

For more information about using the operator panel service functions, see the online help (?).

## Viewing or modifying the status of a migrating partition

Whether the partition is migrating to this system or to another system, you can view the status, stop, or recover the migration from the Integrated Virtualization Manager.

Before you start, complete the following tasks:

1. Ensure that the Integrated Virtualization Manager is at version 1.5 or later. To view the version of the Integrated Virtualization Manager, see “Viewing and updating the code level of the Integrated Virtualization Manager management partition” on page 199.
2. Ensure that the PowerVM Enterprise Edition (or Advanced POWER Virtualization Enterprise Edition) is enabled. For instructions, see “Entering the activation code for PowerVM Editions with the Integrated Virtualization Manager” on page 10.

To view the status of a migrating partition, complete the following steps:

1. From the View/Modify Partitions panel, select the client logical partition that you want to migrate, and click **Status**.
2. To view the status of a migrating partition, review the information on the Migration Status panel.

3. To stop the migration, click **Stop Migration**. When you stop the migration, the Integrated Virtualization Manager (from which the migration was initiated) attempts to reverse all changes and return the migrating partition to the state it was in before the migration began.
4. To recover a migration, click **Recover Migration**. You might need to recover a migration if there is a communication loss between the platform managers, however situation is rare.

### Viewing partition reference codes

Use the Integrated Virtualization Manager to display reference codes for the logical partitions on your managed system. Reference codes provide general system diagnostic, troubleshooting, and debugging information.

To view partition reference codes, do the following:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition for which you want to view reference codes.
3. From the **Tasks** menu, click **Reference Codes**. The Partition Reference Codes panel is displayed.
4. To view a history of reference codes, enter the number of reference codes that you want to view in the **View history** field, and then click **Go**. The panel displays the number of the most recent reference codes that you specified, with the date and time at which each reference code was received.
5. To view the details of a specific reference code, select the option next to the desired reference code. Details about the reference code you selected are displayed in the **Details** area.
6. Click **OK** to close the panel.

## Managing storage devices using the Integrated Virtualization Manager

Use the Integrated Virtualization Manager for storage-management tasks to manage the storage capability of the managed system.

A single storage pool is created automatically when you install the Virtual I/O Server. This storage pool, which is called rootvg, is the *default storage pool*. You might want to consider creating a storage pool in addition to the default rootvg storage pool, and then assign the new storage pool as the default. You can then add more physical volumes to the default storage pool, create virtual disks from the default storage pool, and assign these virtual disks to other logical partitions.

### Creating virtual optical devices using the Integrated Virtualization Manager

You can add a new virtual optical device and mount media to the new device using the Integrated Virtualization Manager.

**Note:** You can also perform this procedure when you use the Create Partition wizard.

Before you start, ensure that the Integrated Virtualization Manager is at version 1.5 or later. To update the Integrated Virtualization Manager, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.

To create a virtual optical device, complete the following steps:

1. From the **Partition Management** menu, click **View/Modify Partitions**.
2. Select the partition for which you want to create a virtual optical device and click **Properties**. The Partition Properties panel is displayed.
3. From the Virtual Optical Devices section, click **Create Device**. A new virtual optical device is created and appears in the table.
4. From the Current Media column of the virtual optical device you just created, click **Modify** to mount media to the new device. The Modify Current Media panel displays.
5. Select the media you want to mount, and click **OK** to return to the Partition Properties panel.

## Modifying virtual disks

You can use the Integrated Virtualization Manager to view the properties of the virtual disks on your managed system, as well as to start virtual disk maintenance tasks.

To view and modify your virtual disks, do the following:

1. From the **Virtual Storage Management** menu, click **View/Modify Virtual Storage**. The View/Modify Virtual Storage panel is displayed.
2. Click the **Virtual Disks** tab. A list of virtual disks is displayed.
3. Select the virtual disk that you want to modify.
4. From the **Tasks** menu, click one of the following:
   - **Properties** to view the properties of the selected virtual disk
   - **Extend** to add storage capacity to the selected virtual disk
   - **Delete** to delete the selected virtual disk and make the storage resources that belonged to that virtual disk available to other virtual disk
   - **Modify partition assignment** to change the virtual disk to which the selected virtual disk is assigned or to set the selected virtual disk so it is not assigned to any logical partition

## Modifying storage pools using the Integrated Virtualization Manager

You can extend a storage pool, reduce or remove a storage pool, and assign a storage pool as the default using the Integrated Virtualization Manager.

To modify storage pools, complete the following tasks:

1. From the **Virtual Storage Management** menu, click **View/Modify Virtual Storage**. The View/Modify Virtual Storage panel is displayed.
2. Click the **Storage Pools** tab. A list of storage pools is displayed.
3. Select the storage pool that you want to modify.
4. From the **Tasks** bar, click one of the following:
   - **Properties** to view the properties of the selected storage pool.
   - **Extend** to add storage to the selected storage pool. To extend logical volume based storage pools, add physical volumes to the storage pool. To extend file based storage pools, add space from the parent storage pool to the file based storage pool.
   - **Reduce** to reduce the size of the selected storage pool. To reduce logical volume based storage pools, remove physical volumes from the storage pool. To reduce the file based storage pool, delete the storage pool.
   - **Assign as default storage pool** to designate the selected storage pool as the default storage pool for this managed system.

## Modifying physical volumes

Use the Integrated Virtualization Manager to view the properties of the physical volumes on your managed system, and to start physical volume maintenance tasks.

A physical volume is an individual logical unit that is identified by a *logical unit number* (LUN). A physical volume can be a hard disk or a logical device on a *storage area network* (SAN). You can either assign a physical volume directly to a logical partition, or you can add a physical volume to a storage pool and create virtual disks from the storage pool.

To view and modify your physical volumes, do the following:

1. From the **Virtual Storage Management** menu, click **View/Modify Virtual Storage**. The View/Modify Virtual Storage panel is displayed.
2. Click the **Physical Volumes** tab to display a list of physical volumes.
3. Select the physical volume that you want to modify.

4. From the **Tasks** menu, click one of the following:
   - **Properties** to view or change the properties of the selected physical volume
   - **Modify partition assignment** to change the logical partition to which the selected physical volume is assigned or to set the physical volume so it is not assigned to any logical partition
   - **Add to storage pool** to add the selected physical volume to the selected storage pool
   - **Remove from storage pool** to remove the selected physical volume from the selected storage pool

### Modifying optical devices using the Integrated Virtualization Manager

You can view and modify physical optical devices and virtual optical media using the Integrated Virtualization Manager.

You can add optical devices to or remove optical devices from any logical partition, whether or not the logical partition is running. If you remove an optical device from a running logical partition, you are prompted to confirm the removal before the optical device is removed.

To modify virtual optical media, the Integrated Virtualization Manager must be at version 1.5 or later. To update the Integrated Virtualization Manager, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.

To modify optical devices, complete the following steps:

1. From the **Virtual Storage Management** menu, click **View/Modify Virtual Storage**. The View/Modify Virtual Storage panel is displayed.
2. Click the **Optical Devices** tab.
3. To change the logical partition assignment setting for a physical optical device, complete the following steps:
   a. From the Physical Optical Devices table, select the optical device that you want to modify.
   b. From the tasks menu, click **Modify partition assignment**. The Modify Optical Device Partition Assignment panel is displayed.
   c. Either change the logical partition to which the optical device is assigned, or set the optical device so it is not assigned to any logical partition, and then click **OK**. The list of optical devices is displayed with the changes you made.
4. To modify virtual optical media, click one of the following tasks from the Virtual Optical Media section:
   - **Extend Library** to extend the size of the media library.
   - **Delete Library** to delete the media library and the files within the library.
   - ***Add Media...** to add an optical media file to the media library and make it available for assignment to a partition.
   - **Modify partition assignment** to change the partition assignment for a media file by changing the virtual optical device to which a media file is assigned. You can assign read-only media to more than one device.
   - **Download** to open or download the selected media file.
   - **Delete** to delete the selected media files from the media library.

## Managing Ethernet using the Integrated Virtualization Manager

Use the Integrated Virtualization Manager for network-management tasks to manage the network connectivity of the managed system.

### Changing the TCP/IP settings on the Virtual I/O Server

Use the Integrated Virtualization Manager to change the TCP/IP settings on the Virtual I/O Server.

Use any role other than the View Only role to perform this task. Users with the View Only role can view the TCP/IP settings, but cannot change them.

Before you can view or modify the TCP/IP settings, you must have an active network interface.

Integrated Virtualization Manager version 1.5.2 supports the use of IPv6 addresses.

**CAUTION:**
**Modifying your TCP/IP settings remotely might result in the loss of access to the current session. Ensure that you have physical console access to the Integrated Virtualization Manager partition prior to making changes to the TCP/IP settings.**

To view or modify the TCP/IP settings, do the following:

1. From the **IVM Management** menu, click **View/Modify TCP/IP Settings**. The View/Modify TCP/IP Settings panel is displayed.
2. Depending on which setting you want to view or modify, select one of the following tabs:
   - **General** to view or modify the host name and the partition communication IP address.

     **Note:** Currently, Integrated Virtualization Manager supports only IPv4 addresses for the partition communication setting.
   - **Network Interfaces** to view or modify the network interface properties, such as the IP address, subnet mask, and the state of the network interface.
   - **Name Services** to view or modify the domain name, name server search order, and domain server search order.
   - **Routing** to view or modify the default gateway.

     **Note:** You can configure both an IPv4 default gateway and an IPv6 default gateway for Integrated Virtualization Manager version 1.5.2.
3. Click **Apply** to activate the new settings.

## Creating a virtual Ethernet adapter

You can create a virtual Ethernet adapter on the management partition and the client partitions using the Integrated Virtualization Manager.

Virtual Ethernet provides Ethernet connectivity between partitions. To create a virtual Ethernet adapter, specify the ID number of the virtual Ethernet network to which you want to connect a corresponding virtual Ethernet adapter that is available for a logical partition. You can also add new adapters or virtual Ethernet network IDs for the partition.

Most tasks associated with IEEE 802.1Q virtual Ethernet are performed using the command-line interface. For detailed command descriptions, see the *Virtual I/O Server and Integrated Virtualization Manager Command Reference*. To view the PDF file of the *Virtual I/O Server and Integrated Virtualization Manager Command Reference* (SA76-0101), approximately 4 MB in size, see sa76-0101.pdf .

To create a virtual Ethernet adapter, complete the following steps:

1. From the **Partition Management** menu, click **View/Modify Partitions**.
2. Select the logical partition to which you want to assign the virtual Ethernet adapter and click **Properties**.
3. Select the **Ethernet** tab.
4. To create a virtual Ethernet adapter on the management partition, complete the following steps:
   a. In the Virtual Ethernet Adapters section, click **Create Adapter**.
   b. Enter the Virtual Ethernet ID and click **OK** to exit the Enter Virtual Ethernet ID window.

   c. Click **OK** to exit the Partition Properties window.
5. To create a virtual Ethernet adapter on a client partition, complete the following steps:
   a. In the Virtual Ethernet Adapters section, select a virtual Ethernet for the adapter and click **OK**.
   b. If no adapters are available, click **Create Adapter** to add a new adapter to the list and then repeat the previous step.

### Viewing virtual Ethernet settings using the Integrated Virtualization Manager

Use the Integrated Virtualization Manager to view the virtual Ethernet settings for the managed system.

Use any role other than View Only to perform the tasks in the **Virtual Ethernet** tab.

To view the virtual Ethernet settings for the managed system, click **View/Modify Virtual Ethernet** from the **I/O Adapter Management**. The **Virtual Ethernet** tab shows information that can be viewed as follows:

- You can view the information by partition, which shows a list of all virtual Ethernets to which each logical partition belongs.
- You can view the information by virtual Ethernet, which shows a list of all logical partitions belonging to each virtual Ethernet.

## Managing system plans

You can create, view, import, export, deploy, and delete system plans on the Integrated Virtualization Manager.

A *system plan* is a specification of the hardware and the logical partitions contained in one or more systems. There are several ways that you can work with system plans. For example, you can import a system plan to the Integrated Virtualization Manager and then deploy the system plan to the managed system. The System Plan Deployment wizard automatically creates logical partitions based on the specifications contained in the system plan. You can also create a system plan based on the current system configuration and then export the system plan to media. Then, you can import the system plan to another system and deploy the system plan to that system.

### Creating a system plan by using the Integrated Virtualization Manager

You can use the Integrated Virtualization Manager to create a new system plan based on an existing system configuration.

When you use the Integrated Virtualization Manager to create a system plan based on the existing managed system, the management partition reads the configuration information on the managed system and stores this information in the system plan.

To create a system plan based on an existing system configuration by using the Integrated Virtualization Manager, complete the following steps:

1. In the navigation area, select **Manage System Plans**. The Manage System Plans page opens.
2. Click **Create/Import system plan** in the toolbar at the top of the System Plans table. The Create/Import System Plan page opens.
3. Select the **Create** option.
4. Enter a System plan file name and plan description for the new system plan.
5. Click **OK**. The Integrated Virtualization Manager generates a new system plan based on the current system configuration and the new system plan appears in the System Plans table.

Now that you have a new system plan, you can export the system plan, import it onto another Integrated Virtualization Manager managed system, and deploy the system plan to that managed system.

**Note:** As an alternative to the Integrated Virtualization Manager Web user interface, you can also use the mksysplan command to accomplish this task.

## Viewing a system plan on the Integrated Virtualization Manager

You can view a system plan on the Integrated Virtualization Manager by using the System Plan Viewer.

The System Plan Viewer uses a navigation tree and tables to display the information in the system-plan file. It includes features such as dynamic table-column sorting and the ability to display EADS boundary lines. The System Plan Viewer is included with the Integrated Virtualization Manager so that it can be accessed from the Integrated Virtualization Manager. However, it requires that you reenter your user ID and password before you can view the system plan.

To view a system plan from the Integrated Virtualization Manager, complete the following steps:

1. From the navigation area, select **Manage System Plans**. The Manage System Plans page opens.
2. Select the system plan that you want to view from the System Plans table.
3. Select **More Tasks** → **View** from the toolbar at the top of the System Plans table. The System Plan Viewer login window opens in a separate browser window.
4. Enter your Integrated Virtualization Manager **Username** and **Password** to log in to the System Plan Viewer.

## Importing a system plan into the Integrated Virtualization Manager

You can import a system-plan file into the Integrated Virtualization Manager management partition. You can then deploy the system plan to the system that the Integrated Virtualization Manager manages.

You must import a system-plan file into the Integrated Virtualization Manager management partition before you can deploy the system plan to the managed system. To deploy the imported system plan successfully, the hardware on the managed system must match or exceed the hardware in the system plan. To import a system-plan file, you must be the prime administrator (padmin). For more information about user roles, refer to User roles.

To import a system-plan file into the Integrated Virtualization Manager management partition, complete the following steps:

1. From the navigation area, select **Manage System Plans**. The Manage System Plans page opens.
2. Click **Create/Import system plan** in the toolbar at the top of the System Plans table. The Create/Import System Plan page opens.
3. Select the **Import** option.
4. Enter the fully qualified path and file name of the system-plan file into the **System plan file name** field. Or, click **Browse** to select the system-plan file from the local file system. The name of the system-plan file must end with the .sysplan file name suffix and can use alphanumeric characters only.
5. Click **OK**. If the Integrated Virtualization Manager returns an error, return to step 4 and verify that the information you entered in this field is correct.

You now can use the imported system-plan file to deploy the system plan to the system that the Integrated Virtualization Manager manages.

**Related concepts**

"User roles" on page 202
Learn about the user roles for the Integrated Virtualization Manager.

## Deploying a system plan by using the Integrated Virtualization Manager

When you deploy a system plan, the Integrated Virtualization Manager creates logical partitions on the managed system according to the specifications in the system plan.

### Requirements for deploying a system plan

When you use a version of the Integrated Virtualization Manager prior to version 1.5.2.0, ensure that the system is in the manufacturing default configuration. More specifically, the system must meet the following requirements:

- Client logical partitions are not configured on the managed system.
- Virtual Ethernet adapters are not configured on the managed system.
- Storage pools are not configured on the managed system.
- Backing devices are not configured on the managed system.
- All of the I/O resources are assigned to the Integrated Virtualization Manager management partition. In the manufacturing default configuration, all the I/O resources are assigned to the Integrated Virtualization Manager. When you add I/O resources to the system, they are assigned to the Integrated Virtualization Manager by default.

When you deploy a system plan with Integrated Virtualization Manager version 1.5.2.0, you can deploy a system plan to a system that is not new or that is not in the manufacturer default configuration. You can deploy a system plan to a system that already has a configured Integrated Virtualization Manager management partition or that has configured client logical partitions. Also, the target system can have any of the previously listed Virtual I/O Server items, such as virtual disks or virtual Ethernet adapters, already configured. However, if the system plan that you intend to deploy contains information about any items that are already configured on the system, the configured items on the target system must exactly match those same items in the system plan. If they do not match exactly, then the system plan either cannot pass validation or the item in the system plan cannot be deployed.

In addition, you no longer must deploy a system plan in its entirety, but can instead partially deploy a system plan on the target system by selecting which logical partitions in the plan to deploy. You then can run the Deploy System Plan Wizard again to deploy the remainder of the logical partitions in the system plan at another time.

In addition to meeting the above requirements based on your version of the Integrated Virtualization Manager, you must meet the following prerequisites:

- The system-plan file exists on the Integrated Virtualization Manager. If the system-plan file does not exist on the Integrated Virtualization Manager, you must import the system-plan file into the Integrated Virtualization Manager.
- The physical hardware is connected and is reporting to the server. If you are deploying a system plan that you created by using the Integrated Virtualization Manager, verify that the hardware and cabling on the target system is identical to that on the source system.
- The physical hardware on the managed system must match exactly to any of the same hardware in the system plan.
- The Integrated Virtualization Manager is not performing any other operations on the managed system.
- You are the prime administrator (padmin). For more information about user roles, refer to User roles.

### Deploying a system plan

To deploy a system plan on a managed system by using the Integrated Virtualization Manager, complete the following steps:

1. In the navigation area of the Integrated Virtualization Manager, select **Manage System Plans**. The Manage System Plans page opens.
2. In the System Plans table, select the system plan that you want to deploy.
3. Select **More Tasks** → **Deploy** from the toolbar at the top of the System Plans table to start the Deploy System Plan Wizard. The System Deployment: Deployment Overview page of the wizard opens.

4. If prompted, choose the managed system to which you want to deploy the system plan and click **Next**. The prompt only occurs if the system plan file contains more than one system. If the system plan does not match the hardware on the managed system to which you want to deploy the plan, the wizard displays a window that informs you of this. Click **OK** to continue or **Cancel** to select a different system plan.
5. Wait for the wizard to validate the managed system and its hardware against the system plan. The validation process can take several minutes.
6. If the validation process completes successfully, click **Next**. If the validation process does not complete successfully, correct the issues indicated by the error messages, click **Cancel** to exit the wizard, and restart this procedure from the beginning. To help you to correct validation issues, you might want to create a system plan that is based on the current configuration of the managed system. Such a system plan can help you to compare the system plan that you want to deploy with the current configuration of the managed system. You can do this by using the Create System Plan task in the Integrated Virtualization Manager, or you can run the following command on the system:

   ```
   mksysplan -f name_of_new_system_plan.sysplan
   ```
7. Review the Deployable Plan Items page, select the logical partitions in the system plan that you want to deploy, and click **Next**. This page indicates the deployable status of the logical partitions that the system plan contains. If a logical partition has a status of partially deployed, the logical partition is selected for deployment automatically, and you cannot change the selection. If a logical partition is deselected for deployment, the wizard does not deploy any dependent entities for that partition, such as backing devices on the Virtual I/O Server.
8. Review the list of deployment plan items on the Deployment page, and click **Finish** to begin deploying the system plan. The Integrated Virtualization Manager creates the specified logical partitions and deploys the specified entities as listed. The deployment process can take several minutes depending on the number of logical partitions and entities to be deployed.

After you finish the deployment of the system plan, complete the following tasks:

- Locate the physical disk I/O adapters that belong to each logical partition and verify that the disk drives that are attached to these physical I/O adapters will support your desired configuration for each logical partition.
- Install operating systems and software on the logical partitions.

**Related concepts**

"User roles" on page 202
Learn about the user roles for the Integrated Virtualization Manager.

**Related tasks**

"Importing a system plan into the Integrated Virtualization Manager" on page 195
You can import a system-plan file into the Integrated Virtualization Manager management partition. You can then deploy the system plan to the system that the Integrated Virtualization Manager manages.

## Exporting a system plan from the Integrated Virtualization Manager

You can export a system-plan file from the Integrated Virtualization Manager and download it to the computer that you use to access the Integrated Virtualization Manager.

To export a system-plan file, you must be the prime administrator (padmin). For more information about user roles, refer to User roles.

To export a system-plan file that is stored on an Integrated Virtualization Manager, complete the following steps:

1. In the navigation area of your Integrated Virtualization Manager, select **Manage System Plans**. The Manage System Plans page opens.
2. Select the system plan that you want to export from the System Plans table

3. Select **More Tasks** → **Export** from the toolbar at the top of the System Plans table. A File Download window for your browser opens.
4. Specify whether to open the system plan by using the browser or to save the system plan to disk and click **OK**. Because you access the system plan by using a browser interface, the browser downloads and exports the system plan. The browser settings control where the system plan downloaded to your local file system.

   **Note:** Your browser might add an extension of .zip to the system plan file name. If this is the case, rename the file to remove the .zip extension to ensure that you can use the system plan file.
   If the Integrated Virtualization Manager returns an error, verify that the information you entered in this window is correct. If necessary, click **Cancel**, return to step 3, and redo the procedure, ensuring that the information you specify at each step is correct.

You can import the system-plan file into a different Integrated Virtualization Manager so that you can deploy the system plan to other managed systems.

**Related concepts**

"User roles" on page 202
Learn about the user roles for the Integrated Virtualization Manager.

### Deleting a system plan from the Integrated Virtualization Manager

You can remove a system plan from the Integrated Virtualization Manager management partition.

Removing a system plan from the Integrated Virtualization Manager management partition does not undo any partition or hardware configuration changes that occurred if the specified system plan was deployed on the managed system.

To remove the system plan from the Integrated Virtualization Manager management partition, complete the following steps:
1. From the navigation area, select **Manage System Plans**. The Managed System Plans page opens.
2. In the System Plans table, select the system plan that you want to delete.
3. Select **More Tasks** → **Remove** from the toolbar at the top of the System Plans table. The Remove System Plans page opens.
4. Confirm that the listed system plan is the one that you want to remove and click **OK** to delete the system plan.

## Updating the Integrated Virtualization Manager

You can update the code level of the management partition and the Virtual I/O Server's firmware microcode using the Integrated Virtualization Manager.

To update the code level of the management partition or the Virtual I/O Server's firmware microcode, complete one of the following procedures:
- Update the current code level of the Integrated Virtualization Manager management partition. For instructions, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.
- Generate a microcode survey of the managed system and download and upgrade the microcode. For instructions, see *Updating the Virtual I/O Server's firmware and device microcode through the Integrated Virtualization Manager with an Internet connection* in the installation guide for your server.
- Update the Virtual I/O Server's firmware and device microcode. For instructions, see *Updating the Virtual I/O Server's firmware and device microcode through the Integrated Virtualization Manager without an Internet connection* in the installation guide for your server.

### Viewing and updating the code level of the Integrated Virtualization Manager management partition

You can view and update the current code level of the Integrated Virtualization Manager management partition.

To update the management partition, complete the following steps:

1. From the **Service Management** menu, click **Updates**.
2. View the current code level of the Integrated Virtualization Manager.
3. Go to the Web site provided on the panel to find the latest available updates and directions for how to apply the updates.

## Installing or replacing a PCI adapter with the system power on in Virtual I/O Server

You might need to install or replace a PCI adapter in the Virtual I/O Server logical partition or in the Integrated Virtualization Manager management partition. Use the procedure in this section to perform this task.

The Virtual I/O Server includes a PCI Hot Plug Manager that is similar to the PCI Hot Plug Manager in the AIX operating system. The PCI Hot Plug Manager allows you to hot plug PCI adapters into the server and then activate them for the logical partition without having to reboot the system. Use the PCI Hot Plug Manager for adding, identifying, or replacing PCI adapters in the system that are currently assigned to the Virtual I/O Server.

### Getting started

**Prerequisites:**

- If you are installing a new adapter, an empty system slot must be assigned to the Virtual I/O Server logical partition. This task can be done through dynamic logical partitioning (DLPAR) operations.
  - If you are using a Hardware Management Console (HMC), you must also update the logical partition profile of the Virtual I/O Server so that the new adapter is configured to the Virtual I/O Server after you restart the system.
  - If you are using the Integrated Virtualization Manager, an empty slot is probably already assigned to the Virtual I/O Server logical partition because all slots are assigned to the Virtual I/O Server by default. You only need to assign an empty slot to the Virtual I/O Server logical partition if you previously assigned all empty slots to other logical partitions.
- If you are installing a new adapter, ensure that you have the software required to support the new adapter and determine whether there are any existing PTF prerequisites to install.
- If you need help determining the PCI slot in which to place a PCI adapter, see PCI adapter placement in the system unit or expansion unit in the PCI Adapter Placement Guide.

Follow these steps to access the Virtual I/O Server, PCI Hot Plug Manager:

1. If you are using the Integrated Virtualization Manager, connect to the command-line interface.
2. Use the **diagmenu** command to open the Virtual I/O Server diagnostic menu. The menus are similar to the AIX diagnostic menus.
3. Select **Task Selection**, then press Enter.
4. At the Task Selection list, select **PCI Hot Plug Manager**.

### Installing a PCI adapter

To install a PCI adapter with the system power on in Virtual I/O Server, do the following:

1. From the PCI Hot Plug Manager, select **Add a PCI Hot Plug Adapter**, then press Enter. The Add a Hot-Plug Adapter window is displayed.

2. Select the appropriate empty PCI slot from those listed, and press Enter. A fast-blinking amber LED located at the back of the server near the adapter indicates that the slot has been identified.
3. Follow the instructions on the screen to install the adapter until the LED for the specified PCI slot is set to the Action state. The adapter installation is performed the same as in a stand-alone AIX logical partition and includes the following sequence of events:
   a. Set the adapter LED to the action state so that the indicator light for the adapter slot flashes
   b. Physically install the adapter
   c. Finish the adapter installation task in **diagmenu**.
4. Run the **cfgdev** command to configure the device for the Virtual I/O Server.

If you are installing a PCI, Fibre Channel adapter, it is now ready to be attached to a SAN and have LUNs assigned to the Virtual I/O Server for virtualization.

## Replacing a PCI Adapter

**Prerequisite:** Before you can remove or replace a storage adapter, you must unconfigure that adapter. See "Unconfiguring storage adapters" on page 115 for instructions.

To replace a PCI adapter with the system power on in Virtual I/O Server, do the following:

1. From the PCI Hot Plug Manager, select **Unconfigure a Device**, then press Enter.
2. Press F4 (or Esc +4) to display the **Device Names** menu.
3. Select the adapter you are removing in the **Device Names** menu.
4. In the **Keep Definition** field, use the Tab key to answer `Yes`. In the **Unconfigure Child Devices** field, use the Tab key again to answer `YES`, then press Enter.
5. Press Enter to verify the information on the **ARE YOU SURE** screen. Successful unconfiguration is indicated by the `OK` message displayed next to the Command field at the top of the screen.
6. Press F4 (or Esc +4) twice to return to the Hot Plug Manager.
7. Select **replace/remove PCI Hot Plug adapter**.
8. Select the slot that has the device to be removed from the system.
9. Select **replace**. A fast-blinking amber LED located at the back of the machine near the adapter indicates that the slot has been identified.
10. Press Enter which places the adapter in the action state, meaning it is ready to be removed from the system.

## Unconfiguring storage adapters

Before you can remove or replace a storage adapter, you must unconfigure that adapter. Storage adapters are generally parent devices to media devices, such as disk drives or tape drives. Removing the parent requires that all attached child devices either be removed or placed in the define state.

Unconfiguring a storage adapter involves the following tasks:

- Closing all applications that are using the adapter you are removing, replacing, or moving
- Unmounting file systems
- Ensuring that all devices connected to the adapter are identified and stopped
- Listing all slots that are currently in use or a slot that is occupied by a specific adapter
- Identifying the adapter's slot location
- Making parent and child devices unavailable
- Making the adapter unavailable

If the adapter supports physical volumes that are in use by a client logical partition, then you might need to perform steps on the client logical partition before unconfiguring the storage adapter. For instructions, see "Preparing the client logical partitions" on page 116. For example, the adapter might be in use because the physical volume was used to create a virtual target device, or it might be part of a volume group used to create a virtual target device.

Follow these steps to unconfigure SCSI, SSA, and Fibre Channel storage adapters:

1. Connect to the Virtual I/O Server command-line interface.
2. Use the oem_setup_env command to close all applications that are using the adapter you are unconfiguring.
3. Type `lsslot-c pci` to list all the hot plug slots in the system unit and display their characteristics.
4. Type `lsdev -C` to list the current state of all the devices in the system unit.
5. Type `unmount` to unmount previously mounted file systems, directories, or files using this adapter.
6. Type `rmdev -l adapter -R` to make the adapter unavailable.

   **Attention:** Do not use the -d flag with the rmdev command for hot plug operations because this action removes your configuration.

## Preparing the client logical partitions

If the virtual target devices of the client logical partitions are not available, the client logical partitions can fail or they might be unable to perform I/O operations for a particular application. If you use the HMC to manage the system, you might have redundant Virtual I/O Server logical partitions, which allow for Virtual I/O Server maintenance and avoid downtime for client logical partitions. If you are replacing an adapter on the Virtual I/O Server and your client logical partition is dependent on one or more of the physical volumes accessed by that adapter, then you might need to take action on the client before you unconfigure the adapter.

The virtual target devices must be in the define state before the Virtual I/O Server adapter can be replaced. Do not remove the virtual devices permanently.

To prepare the client logical partitions so that you can unconfigure an adapter, complete the following steps depending on your situation.

*Table 50. Situations and steps for preparing the client logical partitions*

| **Situation** | **Steps** |
|---|---|
| You have redundant hardware on the Virtual I/O Server for the adapter. | No action is required on the client logical partition. |
| HMC-managed systems only: You have redundant Virtual I/O Server logical partitions that, in conjunction with virtual client adapters, provide multiple paths to the physical volume on the client logical partition. | No action is required on the client logical partition. However, path errors might be logged on the client logical partition. |
| HMC-managed systems only: You have redundant Virtual I/O Server logical partitions that, in conjunction with virtual client adapters, provide multiple physical volumes that are used to mirror a volume group. | See the procedures for your client operating system. |
| You do not have redundant Virtual I/O Server logical partitions. | Shut down the client logical partition.<br>For instructions, see the following topics about shutting down logical partitions:<br>• For systems that are managed by the HMC, see "Shutting down AIX logical partitions using the HMC", and "Shutting down Linux logical partitions using the HMC" in the *Logical Partitioning Guide*.[1]<br>• For systems that are managed by the Integrated Virtualization Manager, see "Shutting down logical partitions" on page 188. |

*Table 50. Situations and steps for preparing the client logical partitions (continued)*

| Situation | Steps |
|---|---|
| **Note:**<br>1. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf . | |

## Creating and modifying user accounts

Use the user-management tasks to manage the Integrated Virtualization Manager user accounts on your managed system.

Use the padmin user account to view, change, or create user accounts.

The management partition on the managed system uses the same user accounts as on the Integrated Virtualization Manager. This means that changes that you make to user accounts using the Integrated Virtualization Manager also apply to the user accounts on the management partition. For example, if you change the password for a user account in the Integrated Virtualization Manager, then you must use the new password when you use that user account to log in to the management partition.

To view a list of Integrated Virtualization Manager user accounts, and to start user maintenance tasks for those user accounts, click **View/Modify User Accounts**.

### User roles

Learn about the user roles for the Integrated Virtualization Manager.

The user role determines which functions a user can access or use. You cannot change the user role that is assigned to a user account after the user account is created. You cannot create user accounts with the same authority as the padmin user account.

The following table lists the user roles available for the Integrated Virtualization Manager.

*Table 51. Integrated Virtualization Manager user roles*

| User role | Description |
|---|---|
| padmin | This role is similar to the root user. Only one padmin user can be created for the Integrated Virtualization Manager. The padmin user account is required to view, change, or create user accounts, and this account can perform all tasks in the Integrated Virtualization Manager. |
| View/Modify | This role is the default type for all users that are not padmin. This role can perform most functions within Integrated Virtualization Manager. The command-line interface calls this role the *Administrator* role. |
| View Only | This role is a read-only role and can perform only list-type (ls) functions. Users with this role do not have the authority to change the system configuration and do not have write permission to their home directories. The command-line interface calls this role the View role. |

*Table 51. Integrated Virtualization Manager user roles (continued)*

| User role | Description |
|---|---|
| Service Representative (SR) | This role allows service representatives to run commands that are required to service the system without being logged in as root. The standard SR login user name is qserv. Some Integrated Virtualization Manager service functions are available only for SR accounts. The service commands for SR accounts include the following:<br>• Run diagnostics, including service aids, such as hot plug tasks, certify, and format.<br>• Run all commands that can be run by a group system.<br>• Configure and unconfigure devices that are not busy.<br>• Use the service aid to update system microcode.<br>• Perform the shutdown and restart operations. |

## Creating user accounts

This topic describes how to create Integrated Virtualization Manager user accounts and set basic properties, such as user ID, password, and role.

Use the padmin user account for this task.

To create a user account, do the following:

1. From the **IVM Management** menu, click **View/Modify User Accounts**. The Create User Accounts panel is displayed.
2. Click ***Create User...**. The Create User Account window displays.
3. Enter the user ID and password, and then confirm the password.
4. Select the appropriate role for the user account, and then click **OK**. The user account is created.

You can create additional user accounts, if necessary. See the online help (?) for more information about user roles.

Only the basic user properties are set up when you create a user account. You can specify additional user properties, such as password restrictions and account expiration date, by changing user properties.

When you create a user account from this panel, the default user role is Administrator. Users with the Administrator user role have authority to perform all tasks except for user maintenance tasks and tasks involving the global command log and the failed login log.

You also cannot create user accounts with the same authority as the padmin user account. The padmin user account can use the Integrated Virtualization Manager to perform all tasks.

## Changing user properties

Use the Integrated Virtualization Manager to change the properties of user accounts, such as number of login retries and the account expiration date.

Use the padmin user account for this task.

To change the properties of a user account, do the following:

1. From the **IVM Management** menu, click **View/Modify User Accounts**. A list of user accounts is displayed.
2. Select the user account for which you want to change the properties.

3. Click **Properties**. The User Properties window displays.
4. On the **User Settings** tab, make the changes you want, and then click **OK**. The list of user accounts is displayed again.

Changes that you make to the settings on the **User Settings** tab take effect the next time that the user logs into the Integrated Virtualization Manager. See the online help (?) for more information about specific user properties.

The management partition on the managed system uses the same user accounts as on the Integrated Virtualization Manager. This means that changes that you make to user accounts using the Integrated Virtualization Manager also apply to management partition user accounts. For example, if you change the password for a user account in the Integrated Virtualization Manager, then you must use the new password when you use that user account to log into the management partition.

## Changing password settings

Learn how to change the password settings and restrictions for Integrated Virtualization Manager user accounts. These settings include the number of weeks until the password expires, minimum password length, and other restrictions.

Use the padmin user account for this task.

To change the password settings for a user account, do the following:

1. From the **IVM Management** menu, click **View/Modify User Accounts**. A list of user accounts is displayed.
2. Select the user account for which you want to change the password settings.
3. Click **Properties**. The User Properties window displays.
4. On the **Password Settings** tab, make the changes that you want, and then click **OK**. The list of user accounts is displayed again.

Changes that you make to the settings on the **Password Settings** tab take effect the next time that the user logs into the Integrated Virtualization Manager. See the online help (?) for more information about specific password settings.

The management partition on the managed system uses the same user accounts as on the Integrated Virtualization Manager. This means that changes that you make to user accounts using the Integrated Virtualization Manager also apply to management partition user accounts. For example, if you change the password for a user account in the Integrated Virtualization Manager, then you must use the new password when you use that user account to log into the management partition.

## Removing user accounts

Learn how to remove Integrated Virtualization Manager user accounts.

Use the padmin user account for this task.

**Attention:** This procedure deletes all user information from the Integrated Virtualization Manager and the management partition. This includes the home directories for those users on the management partition and all files within those directories. To preserve the files within the home directories, use the command-line interface on the management partition to copy the files to another location before removing the user accounts.

To remove a user account, do the following:

1. From the **IVM Management** menu, click **View/Modify User Accounts**. A list of user accounts is displayed.

2. Select the user account that you want to remove.
3. Click **Remove account**. The Remove User Accounts window displays, which lists the user accounts you selected to remove.
4. Click **OK** to remove the user account. The list of user accounts is displayed again, and the user account you removed is no longer displayed.

You can select multiple user accounts to remove. For more information about removing user accounts, see the online help (?).

## Changing user passwords

Learn how to change user passwords in the Integrated Virtualization Manager.

Use the padmin user account for this task.

To change a user password, do the following:
1. From the **IVM Management** menu, click **View/Modify User Accounts**. A list of user accounts is displayed.
2. Select the user account for which you want to change the password.
3. Click **Change password**. The Change Password window displays.
4. Enter the new password.
5. Confirm the new password, and then click **OK**. The password is changed, and the list of user accounts is displayed again.

The next time that the user logs in to Integrated Virtualization Manager, the password change takes effect, and the user is required to change it.

The management partition on the managed system uses the same user accounts as the Integrated Virtualization Manager. This means that the password change that you make here also applies to the management partition user account.

Users can change their own user passwords by clicking **Edit my profile** in the toolbar.

## Editing your user profile

Use the Integrated Virtualization Manager to edit your user profile. Specifically, learn how to change your user password.

You must be logged in with the user account for which you want to change the password.

To change the password for your user account, do the following:
1. From the toolbar, click **Edit my profile**. The **Edit My Profile** dialog box is displayed.
2. Type the current password, and then type the new password.
3. Confirm the new password, and then click **OK**. The password is changed, and the Integrated Virtualization Manager page is displayed.

The password change takes effect the next time that you log into the Integrated Virtualization Manager.

The management partition on the managed system uses the same user accounts as on the Integrated Virtualization Manager. This means that the password change that you make here also applies to the management partition user account.

The padmin user account can change passwords for any user account.

## Troubleshooting the Integrated Virtualization Manager

Use service-management tasks to maintain and troubleshoot the Integrated Virtualization Manager.

Use the service management tasks to maintain your managed system so that it is running and up to date.

## Using Service Focal Point for the Integrated Virtualization Manager

Learn about using the Service Focal Point for the Integrated Virtualization Manager to help you manage problems on your system.

Service Focal Point for the Integrated Virtualization Manager is an application that allows you to manage serviceable events, create serviceable events, manage dumps, and collect vital product data (VPD).

## Backing up and restoring partition data

Use the Integrated Virtualization Manager to back up or restore the partition configuration information on your managed system. You can download an existing backup of the partition configuration, generate a new backup, upload a saved backup, or restore the existing backup.

To back up or restore partition data, complete the following steps:

1. From the **Service Management** menu, click **Backup/Restore**. The Backup/Restore page is displayed, which includes the **Partition Configuration Backup/Restore** tab, the **Management Partition Backup/Restore** tab, and the **File and Virtual Media Backup/Restore** tab.
2. To download an existing backup of the partition configuration, generate a new backup, upload a saved backup, or restore the existing backup, click the **Partition Configuration Backup/Restore** tab.
3. To view instructions for backing up and restoring the data on your management partition using the **backupios** command, click the **Management Partition Backup/Restore** tab.

For more information about specific tasks for backing up and restoring partition data, see the online help (?).

You can use the Integrated Virtualization Manager version 1.5.1.1 to back up and restore virtual optical media files and files in your user /home directory. For more information about how to do this, see “Backing up virtual media and user files to tape” and “Restoring virtual media and user files from tape” on page 207.

## Backing up virtual media and user files to tape

Use the Integrated Virtualization Manager to back up files in your user /home directory and virtual media files from your managed system to tape.

You must have a tape device mounted on the managed system to complete this task.

To back up user files or virtual media files to tape, complete the following steps:

1. From the **Service Management** menu, click **Backup/Restore**. The Backup/Restore page is displayed.
2. Click the **File and Virtual Media Backup/Restore** tab.
3. In the **Managed System File** table, select the files that you want to back up to tape. The /home/padmin directory is listed as a single entry.

   Click **[+] Show Files** to have the table list all the files within the directory for individual selection. Click **[-] Hide Files** to have the table list only the /home/padmin directory.

   By selecting the directory entry you can back up all the files in the directory by default.
4. Click **Generate Command**. Integrated Virtualization Manager updates the page by replacing the **Managed System File** table with an informational message that contains the command that you need to run to back up the selected files.

5. Copy the command that the Integrated Virtualization Manager generated and open a terminal session window.
6. Paste the command into the terminal window and run it to back up the selected files to a tape device.

You also can use Integrated Virtualization Manager to restore files in your user /home directory and virtual media files from tape. For more information see “Restoring virtual media and user files from tape.”

## Restoring virtual media and user files from tape

Use the Integrated Virtualization Manager to restore files in your user /home directory and virtual media files from tape to your managed system.

You must have a tape device mounted on the managed system to complete this task.

To restore user files or virtual media files from tape, complete the following steps:

1. From the **Service Management** menu, click **Backup/Restore**. The Backup/Restore page is displayed.
2. Click the **File and Virtual Media Backup/Restore** tab.
3. Click **List Tape Contents** to view a list of all files on the specified tape device. When the process finishes reading the tape, you can view the list of files in the **Tape Device File** table.
4. In the **Tape Device File** table, select the files that you want to restore to the managed system from tape.
5. Click **Generate Command**. Integrated Virtualization Manager updates the page by replacing the **Tape Device File** table with an informational message that contains the command that you need to run to restore the selected files.
6. Copy the command that the Integrated Virtualization Manager generated and open a terminal session window.
7. Paste the command into the terminal window and run it to restore the selected files to the managed system. The command only restores files to those directories to which your user ID has write-access authority. If you have selected to restore a file to a directory to which you do not have such authority, the command cannot restore that particular file.

You also can use Integrated Virtualization Manager to back up files in your user /home directory and virtual media files on the managed system to tape. For more information see “Backing up virtual media and user files to tape” on page 206.

## Viewing application logs

View the application log entries on your managed system. *Application logs* are files that contain events and errors generated by the Integrated Virtualization Manager.

To view the application logs, do the following:

1. From the **Service Management** menu, click **Application Logs**. The Application Logs panel is displayed.
2. To modify the selection criteria, select the desired filters, and then click **Apply**. Click **Reset** to reset the filter information to the default values.

For more information about the selection criteria and filters, see the online help (?).

## Viewing application log properties

Use the Integrated Virtualization Manager to view the properties of the application log entries on your managed system.

To view the properties of the application logs, do the following:

1. From the **Service Management** menu, click **Application Logs**. The Application Logs panel is displayed.
2. Select the application log for which you want to view the properties.
3. From the **Tasks** menu, click **Properties**. The **Log Properties** dialog box is displayed.
4. Click **OK** or **Cancel** to close the dialog box. The Application Logs panel is displayed.

For more information about the specific properties of the application logs, see the online help (?).

## Monitoring tasks

View and monitor the most recent 40 tasks that are running on the Integrated Virtualization Manager.

To view the properties of the tasks, do the following:

1. From the **Service Management** menu, click **Monitor Tasks**. The Monitor Tasks panel is displayed.
2. Select the task for which you want to view the properties.
3. Click **Properties**. The Task Properties dialog box is displayed.
4. Click **Cancel** to close the dialog box. The Monitor Tasks panel is displayed.

## Viewing hardware inventory

Use the Integrated Virtualization Manager to list the devices on your managed system, including device name, status, type of device, and physical location code.

To list the devices on your managed system, do the following:

1. From the **Service Management** menu, click **Hardware Inventory**. The Hardware Inventory panel is displayed, which includes a list of hardware devices.
2. To sort the list by any of the categories, such as device name or status, click the appropriate header.

This list includes any device with a device name, including both physical devices and virtual devices. Using this page is equivalent to using the **lsdev** command in the command-line interface.

For more information about the categories, see the online help (?).

# Connecting an HMC to a system managed by the Integrated Virtualization Manager

Learn how to connect an servers system that is managed by the Integrated Virtualization Manager to become an servers system managed by a Hardware Management Console (HMC).

Connecting an HMC to a system that is managed by the Integrated Virtualization Manager automatically disables the Integrated Virtualization Manager. The HMC assumes management control of the system. Because the management of the system is changed, you must re-create your logical partition configuration either manually or from backups.

To change the management of a system from the Integrated Virtualization Manager to an HMC, do the following:

1. Create a backup of your partition configuration using the Integrated Virtualization Manager, and download it to your local system. For instructions, see “Backing up and restoring partition data” on page 206. You can use the backup text file as a reference for re-creating your partitions in step 4.

2. Connect the HMC to the system. For instructions, see the *Installation and Configuration Guide for the Hardware Management Console*. To view the PDF file of the *Installation and Configuration Guide for the Hardware Management Console* (SA76-0084), approximately 3 MB in size, see sa76-0084.pdf. The managed system is in Recovery state on the HMC.
3. Initialize the profile data using the HMC interface. For instructions, see the *Operations Guide for the Hardware Management Console and Managed Systems*. To view the PDF file of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), approximately 4 MB in size, see sa76-0085.pdf. This action clears the partition configuration.
4. Using the backup text file that you created in step 1, re-create your partitions using the HMC. For instructions, see the *Operations Guide for the Hardware Management Console and Managed Systems*. To view the PDF file of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), approximately 4 MB in size, see sa76-0085.pdf.

# Chapter 5. Live Partition Mobility

Live Partition Mobility, a component of the PowerVM Enterprise Edition (or Advanced POWER Virtualization Enterprise Edition) hardware feature, provides the ability to move AIX and Linux logical partitions from one system to another. The migration procedure transfers the system environment including the processor state, memory, attached virtual devices, and connected users. Active Partition Mobility and inactive Partition Mobility are migration types available to users of POWER6 processor-based servers.

*Active Partition Mobility* allows you to move AIX and Linux running logical partitions, including its operating system and applications, from one system to another. The logical partition and the applications running on that migrated logical partition do not need to be shut down. This type of migration allows you to balance workloads and resources among servers without any impact to your users.

*Inactive Partition Mobility* allows you to move a powered off AIX and Linux logical partition from one system to another. Inactive Partition Mobility is a reliable method to move a logical partition with minimal interaction from the system administrator.

## Concepts for Partition Mobility

Before you start moving logical partitions from one server to another, it is essential that you understand the concepts behind this function.

The purpose of this information is to familiarize you with the configuration required for Partition Mobility.

### Benefits of Partition Mobility

Learn about the advantages and applications of Partition Mobility.

Partition Mobility provides systems management flexibility and is designed to improve system availability. For example:

- You can avoid planned outages for hardware or firmware maintenance by moving logical partitions to another server and then performing the maintenance. Partition Mobility can help because you can use it to work around scheduled maintenance activities.
- You can avoid downtime for a server upgrade by moving logical partitions to another server and then performing the upgrade. This allows you to continue your work without disruption.
- If a server indicates a potential failure, you can move its logical partitions to another server before the failure occurs. Partition Mobility can help avoid unplanned downtime.
- You can consolidate workloads running on several small, underused servers onto a single large server.
- You can move workloads from server to server to optimize resource use and workload performance within your computing environment. With active Partition Mobility, you can manage workloads with minimal downtime.

However, while Partition Mobility provides many benefits, it does not do the following:

- Partition Mobility does not provide automatic workload balancing.
- Partition Mobility does not provide a bridge to new functions. Logical partitions must be restarted and possibly reinstalled to take advantage of new features.

### Active Partition Mobility

Learn more about the process and performance considerations of active Partition Mobility.

Active Partition Mobility is the ability to move AIX and Linux running logical partitions, including its operating system and applications, from one server to another without disrupting the operation of that logical partition. The mobile partition retains its name, partition profiles, and current configuration. A logical partition must have the following characteristics to be ready for active Partition Mobility:

- The logical partition is in the Running state.
- The logical partition has a unique name that is not on the destination server.
- The destination Virtual I/O Server logical partition must have unused virtual slots available.
- The logical partition has only the two default virtual serial I/O adapters assigned to its current configuration.
- All virtual SCSI disks must be mapped to logical unit numbers (LUN)s visible on external Storage Area Network (SAN) and accessible to the Virtual I/O Server on the destination server.
- The logical partition must not have physical adapters or a Host Ethernet Adapter. (Host Ethernet Adapter is sometimes referred to as Integrated Virtual Ethernet.)
- The logical partition is not a Virtual I/O Server logical partition.
- The logical partition is not part of a partition workload group.
- The logical partition is not used as the redundant error path reporting logical partition.
- The logical partition is not using huge pages.
- The logical partition is not using barrier synchronization register (BSR).
- The logical partition does not have any of its virtual SCSI disks defined as logical volumes in any Virtual I/O Server.

## How active Partition Mobility works on the Hardware Management Console and the Integrated Virtualization Manager

Learn more about the end-to-end process of active Partition Mobility.

Active Partition Mobility lets you move a running logical partition, including its operating system and applications, from one server to another without disrupting the operation of that logical partition. Following are the steps that take place during this process:

1. The user ensures that all requirements are satisfied and all preparation tasks are completed.
2. The user initiates active Partition Mobility using the Partition Migration wizard on the Hardware Management Console (HMC) or starts the migration task on the Integrated Virtualization Manager.
3. The HMC or the Integrated Virtualization Manager verifies the Partition Mobility environment.
4. The HMC or the Integrated Virtualization Manager prepares the source and destination environments for active Partition Mobility.
5. The HMC or the Integrated Virtualization Manager transfers the logical partition state from the source environment to the destination environment. This includes all the partition profiles associated with the mobile partition.
   - The source mover service partition extracts the logical partition state information from the source server and sends it to the destination mover service partition over the network.
   - The destination mover service partition receives the logical partition state information and installs it on the destination server.
6. The HMC or the Integrated Virtualization Manager suspends the mobile partition on the source server. The source mover service partition continues to transfer the logical partition state information to the destination mover service partition.
7. The hypervisor resumes the mobile partition on the destination server.
8. The HMC or the Integrated Virtualization Manager completes the migration. All resources that were consumed by the mobile partition on the source server are reclaimed by the source server, including:
   - The source Virtual I/O Server unlocked, unconfigured, or undefined virtual resources on the source servers.

- The HMC or the Integrated Virtualization Manager removes the hosting virtual adapter slots from the source Virtual I/O Server partition profiles as required.

9. The user performs postrequisite tasks, such as:
   - Adding the mobile partition to a partition workload group
   - Adding dedicated I/O adapters

**System characteristics that will not change after a partition migration:** After performing a partition migration, some logical partition attributes will remain the same on the destination server. The following logical partition attributes will not change:

- The logical partition name
- The logical partition type (dedicated processor or shared processor)
- The logical partition configuration
- The processor architecture
- The Simultaneous Multi-Threading (SMT) state of each processor
- The virtual MAC addresses, IP addresses, and LUN mapping to the target devices

**System characteristics that might change after a logical partition migration:** Some system attribute information might change after a logical partition migration. The following system characteristics might change:

- The logical partition ID number
- The machine type, model, and serial number
- The model class of the underlying server
- The processor version and type
- The processor frequency
- The affinity characteristics of the logical memory blocks (LMB)
- The maximum number of hot pluggable and installed physical processors
- The L1 and L2 cache size

## Validating active Partition Mobility

Use this information to learn how the Hardware Management Console (HMC) Partition Migration wizard and the Integrated Virtualization Manager migration function validates your system configuration.

Before you attempt to migrate an active logical partition, you need to validate your environment. You can use the validation function on the HMC or Integrated Virtualization Manager to validate your system configuration. If the HMC or Integrated Virtualization Manager detects a configuration or connection problem, it displays an error message with information to help you resolve the problem. The validation function on the HMC and Integrated Virtualization Manager performs the following tasks:

- Checks the source and destination systems, POWER6 hypervisor, Virtual I/O Servers, and mover service logical partitions for active partition migration capability and compatibility.
- Checks that the Resource Monitoring and Control (RMC) connections to the mobile partition, the source and destination Virtual I/O Servers, and the connection between the source and destination mover service partitions are established.
- Checks that there are no physical adapters in the mobile partition and that there are no virtual serial adapters in virtual slots higher than 1.
- Checks that no client virtual SCSI disks on the mobile partition are backed by logical volumes and that no disks map to internal disks.
- Checks the mobile partition, its operating system, and its applications for active migration capability.
- Checks that the logical memory block size is the same on the source and destination servers.
- Ensures that the operating system on the mobile partition is AIX or Linux.
- Ensures that the logical partition is not the redundant error path reporting logical partition.

- Ensures that the service logical partition is not configured with barrier synchronization registers (BSR).
- Ensures that the mobile partition is not configured with huge pages.
- Ensures that the mobile partition does not have a Host Ethernet Adapter (or Integrated Virtual Ethernet).
- Checks that the logical partition state is Active or Running.
- Checks that the mobile partition is not in a partition workload group.
- Checks the uniqueness of the mobile partition's virtual MAC addresses.
- Checks that the required Virtual LAN IDs are available on the destination Virtual I/O Server.
- Checks that the mobile partition's name is not already in use on the destination server.
- Checks the number of current active migrations against the number of supported active migrations.
- Check that the necessary resources (processors, memory, and virtual slots) are available to create a shell logical partition on the destination system with the exact configuration of the mobile partition.
- Ensures that the operating system in the mobile partition is capable of being migrated. AIX passes the check-migrate request to those applications and kernel extensions that have registered to be notified of dynamic reconfiguration events. The operating system either accepts or rejects the migration.

## Inactive Partition Mobility

Learn more about the end-to-end process of inactive Partition Mobility.

Inactive Partition Mobility lets you move a powered off AIX and Linux logical partition, from one server to another. Because the HMC always migrates the latest activated profile, an inactive logical partition that has never been activated cannot be migrated. The Integrated Virtualization Manager allows you to migrate a logical partition that has never been activated. During the inactive migration, the mobile partition maintains its name and its inactive state. Its virtual I/O resources are assigned and remapped to the appropriate Virtual I/O Server logical partitions on the destination system. Its processor and memory resources remain unassigned until you activate the logical partition on the HMC.

A logical partition must have the following characteristics to be ready for inactive Partition Mobility:

- The logical partition is in a Not Activated state.
- The logical partition can use huge pages.
- The logical partition can use barrier synchronization registers (BSR).
- The logical partition must not have any physical adapters. The HMC and the Integrated Virtualization Manager will automatically remove physical adapters.
- The destination server must have enough processors and memory to support the mobile partition's configuration.
- The logical partition is not a Virtual I/O Server logical partition.
- The logical partition is not part of a partition workload group.
- The logical partition is not used as the redundant error path reporting logical partition.
- The logical partition has only the two default virtual serial I/O adapters assigned to its current configuration.
- The logical partition does not have any of its virtual SCSI disks defined as logical volumes in any Virtual I/O Server .
- All virtual SCSI disks must be mapped to LUNs that are visible on external SAN storage.

### How inactive Partition Mobility works

Learn more about the end-to-end process of inactive Partition Mobility.

With inactive Partition Mobility, you can move a logical partition that is powered off from one server to another.

1. The user ensures that all requirements are satisfied and all preparation tasks are completed.

2. The user shuts down the mobile partition.
3. The user initiates inactive Partition Mobility using the Partition Migration wizard on the HMC or initiates the migration task using the Integrated Virtualization Manager.
4. The HMC and the Integrated Virtualization Manager verifies the Partition Mobility environment.
5. The HMC and the Integrated Virtualization Manager prepares the source and destination environments for inactive Partition Mobility.
6. The HMC and the Integrated Virtualization Manager transfers the partition state from source environment to the destination environment. This includes all the partition profiles associated with the mobile partition on the HMC.
7. The HMC and the Integrated Virtualization Manager completes the migration. This means that all resources that were consumed by the mobile partition on the source server are reclaimed by the source server, including:
   - The source Virtual I/O Server's unlocked, unconfigure, or undefine virtual resources on the source and destination servers.
   - The HMC removes the hosting virtual adapter slots from the source HMC partition profiles.
8. The user activates the mobile partition on the destination server.
9. The user performs postrequisite tasks, such as:
   - Establishing virtual terminal connections
   - Adding the mobile partition to a partition workload group

**System characteristics that will not change after a partition migration:** After performing a partition migration, some logical partition attributes will remain the same on the destination server. The following logical partition attributes will not change:

- The logical partition name
- The logical partition type (dedicated processor or shared processor)
- The logical partition configuration
- The processor architecture
- The Simultaneous Multi-Threading (SMT) state of each processor
- The virtual MAC addresses, IP addresses, and LUN mapping to the target devices

**System characteristics that might change after a logical partition migration:** Some system attribute information might change after a logical partition migration. The following system characteristics might change:

- The logical partition ID number
- The machine type, model, and serial number
- The model class of the underlying server
- The processor version and type
- The processor frequency
- The affinity characteristics of the logical memory blocks (LMB)
- The maximum number of hot pluggable and installed physical processors
- The L1 and L2 cache size

## Validating inactive Partition Mobility

Learn how the HMC and the Integrated Virtualization Manager validate your system configuration.

Before you attempt to migrate an inactive logical partition, you need to validate your environment. You can use the validation function on the HMC or Integrated Virtualization Manager to validate your system configuration. If the HMC or Integrated Virtualization Manager detects a configuration or connection problem, it displays an error message with information to help you resolve the problem. The validation function on the HMC and Integrated Virtualization Manager performs the following tasks:

- Checks the Virtual I/O Server and POWER6 hypervisor migration capability and compatibility on the source and destination.
- Checks that resources (processors, memory, and virtual slots) are available to create a shell logical partition on the destination system with the exact configuration of the mobile partition.
- Verifies the Resource Monitoring and Control (RMC) connections to the source and destination Virtual I/O Servers.
- Ensures that the logical partition name is not already in use at the destination server.
- Checks for virtual MAC address uniqueness.
- Checks that the required Virtual LAN IDs are available on the destination Virtual I/O Server.
- Checks that the logical partition is in the Not Activated state.
- Ensures that the mobile partition is an AIX or Linux logical partition.
- Ensures the mobile partition is not the redundant error path reporting logical partition or a service logical partition.
- Ensure the mobile partition is not a member of a partition workload group.
- Ensures that the mobile partition has an active profile on the HMC.
- Checks the number of current inactive migrations against the number of supported inactive migrations.
- Checks that all required I/O devices are connected to the mobile partition through a Virtual I/O Server, that is, there are no physical adapters.
- Verifies that the virtual SCSI disks assigned to the logical partition are accessible by the Virtual I/O Servers on the destination system.
- Creates the virtual adapter migration map that associates adapters on the source Virtual I/O Servers with adapters on the destination Virtual I/O Servers.
- Ensures that no virtual SCSI disks are backed by logical volumes and that no virtual SCSI disks are attached to internal disks (not on the SAN).

# Using the HMC for Live Partition Mobility

Learn more about using the Hardware Management Console to migrate an active or inactive logical partition.

## HMC environment

Use this information to help gain an understanding of an active or inactive partition migration using the Hardware Management Console.

### Source and destination servers

Learn how to set up the environment for the source and destination servers before you migrate a logical partition.

Two servers are involved in Partition Mobility. The *source server* is the server from which you want to move the logical partition, and the *destination server* is the server to which you want to move the logical partition. The source and destination servers must be POWER6 processor-based servers to participate in Partition Mobility. The destination server must have enough available processor and memory resources to allow the mobile partition to run on its server.

#### Huge pages

Huge pages can improve performance in specific environments that require a high degree of parallelism, such as in DB2® partitioned database environments. You can specify the minimum, desired, and maximum number of huge pages to assign to a logical partition when you create the logical partition or partition profile.

A logical partition cannot participate in active Partition Mobility if huge pages are used. However, an inactive partition migration can be performed if the mobile partition uses huge pages. The partition profile will maintain the huge page resources, but the specified number of huge page resources may not be available on the destination server, in which case the logical partition will boot without some or all of these huge pages after the inactive migration.

### Barrier synchronization register (BSR)

The barrier synchronization register (BSR) is a memory register that is located on certain processors based on POWER technology. A parallel-processing application running on AIX can use a BSR to perform barrier synchronization, which is a method for synchronizing the threads in the parallel-processing application.

A logical partition cannot participate in active partition migration if BSR is used. However, you can use inactive Partition Mobility if you do not want to disable BSR.

## Hardware Management Console

The *HMC* is a system that controls managed systems, including the management of logical partitions and use of Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information for analysis. You can use the *HMC* to configure and control one or more managed systems. The HMC must be at version 7 or later to participate in Live Partition Mobility. During the migration, there is one HMC or redundant HMC pair involved that manages both the source and destination servers. The HMC can handle multiple migrations simultaneously. However, the maximum number of concurrent partition migrations is limited by the processing capacity of the HMC.

The Partition Migration wizard that is provided on the HMC helps you validate and complete a partition migration. The HMC determines the appropriate type of migration to use based on the state of the logical partition. If the logical partition is in the *Running* state, then the migration is active. If the logical partition is in the *Not Activated* state, then the migration is inactive. Before the migration starts, the HMC validates your logical partition environment. During this validation check, the HMC can determine if your migration will be successful. If your migration validation fails, the graphical user interface provides error messages and suggestions to help you resolve your configuration problem.

## Source and destination Virtual I/O Server logical partitions

Learn more about the Virtual I/O Server component of the Partition Mobility environment.

Partition Mobility requires a Virtual I/O Server logical partition on the source server and a Virtual I/O Server logical partition on the destination server. Both Virtual I/O Server logical partitions must provide storage and networking resources to the mobile partition. This allows the mobile partition access to the same storage from both the source and destination servers.

For active Partition Mobility, both the source and the destination Virtual I/O Server logical partitions must be designated as mover service partitions. A *mover service partition* is a Virtual I/O Server logical partition with the following characteristics:

- The mover service partition attribute indicates that the Virtual I/O Server logical partition is capable of supporting active partition migration.
- Both Virtual I/O Servers must be at version 1.5 or later.

The source and destination mover service partitions communicate with each other over the network. On both the source and destination servers, the Virtual Asynchronous Services Interface (VASI) device provides communication between the mover service partition and the hypervisor. These connections facilitate active Partition Mobility as follows:

- On the source server, the mover service partition extracts the logical partition state information of the mobile partition from the hypervisor.

- The mover service partition on the source server sends the logical partition state information to the mover service partition on the destination server.
- On the destination server, the mover service partition installs the logical partition state information on the hypervisor.

## Mobile partition

Learn more about how Partition Mobility impacts the mobile partition configuration.

A *mobile partition* is a logical partition that you want to move from the source server to the destination server. You can migrate a running mobile partition or a powered off mobile partition from the source server to the destination server.

**Note:** Before you attempt the migration, verify that the mobile partition is able to migrate to the destination server.
The HMC creates a migration profile for the mobile partition on the destination server that matches the current configuration of the logical partition. During the migration, the HMC migrates all of the profiles associated with the mobile partition to the destination server. Only the current partition profile (or a new one, if specified) will be converted during the migration process. This conversion will include the mapping of the client virtual SCSI slot to the corresponding target virtual SCSI slot on the destination Virtual I/O Server, if required.

A logical partition cannot be migrated if any logical partition exists on the destination server with the same name. The HMC creates a new migration profile containing the logical partition's current state if you do not specify a profile name. The profile replaces the existing profile that was last used to activate the logical partition. If you specify an existing profile name, the HMC replaces that profile with the new migration profile. If you want to keep the logical partition's existing profiles, specify a new and unique profile name before the migration begins.

**Considerations for configuring I/O for the mobile partition:**

Do not assign any physical or required I/O adapters to a mobile partition using the active partition migration. All of the I/O adapters on the mobile partition must be virtual devices. To remove the physical adapters on the mobile partition, you can use the dynamic logical partition removal task.

A mobile partition with dedicated adapters can participate in inactive Partition Mobility; however, the dedicated adapters will be removed from the partition profile. Thus, the logical partition will boot with only virtual I/O resources after an inactive migration. If dedicated I/O resources were assigned to the logical partition on the source server, these resources will become available when the logical partition is deleted from the source server.

## Networking

In Partition Mobility, the network between the two mover service partitions is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

The virtual LAN must be bridged to a physical network using a Shared Ethernet Adapter in the Virtual I/O Server logical partition. Your LAN must be configured so that the mobile partition can continue to communicate with other necessary clients and servers after a migration is completed.

Active Partition Mobility has no specific requirements on the mobile partition's memory size or the type of network that is connecting the mover service partitions. The memory transfer is a procedure that does not interrupt a mobile partition's activity and may take time when a large memory configuration is busy on a slow network. Because of this, you might want to use a high-bandwidth connection, such as Gigabit Ethernet or faster, between the Virtual I/O Server logical partitions that are providing the mover service partition capability.

The maximum distance between the source and destination systems is dictated by the network and storage configuration used by the systems, the ability of the applications to continue to operate when its storage is separated from the server by such a distance, and the requirement that the source and destination systems must be managed by the same HMC. If both systems are on the same network, connected to the same shared storage, and managed by the same HMC, then active Partition Mobility validation will succeed. The time it takes to move the logical partition and the application performance after a move across a long distance is dependent on the effective network distance between the source and destination systems and application sensitivity to increased storage latency.

### Storage configuration for Partition Mobility

Learn more about storage configuration requirements for Partition Mobility.

The mobile partition moves from one server to another by the source server that is sending the logical partition state information to the destination server over a local area network (LAN). However, partition disk data cannot pass from one system to another system over a network. Thus, for Partition Mobility to succeed, the mobile partition must use storage resources virtualized by a storage area network (SAN) so that it can access the same storage from both the source and destination servers.

### Software applications that recognize migrations

Software applications might be designed to recognize and adapt to changes in the system hardware after being moved from one system to another.

Most software applications running in AIX and Linux logical partitions will not require any changes to work correctly during active Partition Mobility. Some applications may have dependencies on characteristics that change between the source and destination servers and other applications might need to adjust to support the migration.

Examples of applications that would benefit if they were Partition Mobility aware:

- Software applications that use processor and memory affinity characteristics to tune their behavior because affinity characteristics may change as a result of migration. The application's functionality remains the same, but performance variations may be observed.
- Applications that use processor binding will maintain their binding to the same logical processors across migrations, but in reality the physical processors will change. Binding is usually done to maintain hot caches, but the physical processor move will require a cache hierarchy on the destination system. This usually occurs very quickly and should not be visible to the users.
- Applications that are tuned for given cache architectures, such as hierarchy, size, line-size, and associativity.
- Performance analysis, capacity planning, and accounting tools and their agents are usually migration-aware because the processor performance counters may change between the source and destination servers, as may the processor type and frequency. Additionally, tools that calculate an aggregate system load based on the sum of the loads in all hosted logical partitions must be aware that a logical partition has left the system or that a new logical partition arrived.
- Workload managers

## Requirements for Partition Mobility using the Hardware Management Console

Learn more about the software and hardware requirements for the Partition Mobility using the Hardware Management Console.

The hardware and software that are required to use Partition Mobility varies depending on whether you are migrating an active or inactive logical partition. Make sure that your Partition Mobility environment meets minimum requirements before you migrate your logical partition.

## HMC requirements

The following table shows the software requirements for the HMC.

*Table 52. HMC software requirements*

| HMC requirement | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The HMC version and release must be at V7R320 or later.<br>To determine the current HMC version and update it if necessary, see the following instructions:<br>• Determining your HMC machine code version and release in the *Operations Guide for the Hardware Management Console and Managed Systems*[1].<br>• Getting HMC machine code fixes and upgrades in the *Operations Guide for the Hardware Management Console and Managed Systems*[1]. | X | X |
| **Note:**<br>1. To view the PDF file of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), approximately 4 MB in size, see sa76-0085.pdf . | | |

## Source and destination server requirements

The following table shows the hardware requirements for the source and destination server.

*Table 53. Source and destination server requirements*

| Server requirement | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The source and destination server must be one of the following POWER6 models:<br>• 03E/4A<br>• 04E/8A<br>• 17M/MA | X | X |

*Table 53. Source and destination server requirements (continued)*

| Server requirement | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The source and destination servers must both be at firmware level 01EX320 or later.<br>To determine the current firmware level and update it if necessary, see the following instructions:<br>• Using the HMC to view the existing firmware (Licensed Internal Code) levels in the *Operations Guide for the Hardware Management Console and Managed Systems*[1].<br>• Getting server firmware and power subsystem firmware fixes and upgrades in the *Operations Guide for the Hardware Management Console and Managed Systems*[1]. | X | X |
| **Note:**<br>1. To view the PDF file of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), approximately 4 MB in size, see sa76-0085.pdf . | | |

## Source and destination Virtual I/O Server logical partition requirements

The following table shows the software requirements for the source and destination Virtual I/O Server logical partition.

*Table 54. Source and destination Virtual I/O Server logical partition software requirements*

| Virtual I/O Server logical partition requirement | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The PowerVM Enterprise Edition (or Advanced POWER Virtualization Enterprise Edition) hardware feature must be purchased and activated to use Partition Mobility. For more information about PowerVM Enterprise Edition, see PowerVM Editions. For instructions about activating the PowerVM Enterprise Edition hardware feature, see Entering the activation code for PowerVM Editions using the HMC version 7. | X | X |
| At least one Virtual I/O Server logical partition must be installed and activated on both the source and destination servers.<br>For instructions, see Installing the Virtual I/O Server. | X | X |

Table 54. Source and destination Virtual I/O Server logical partition software requirements (continued)

| Virtual I/O Server logical partition requirement | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The source and destination Virtual I/O Server logical partitions must be at release level 1.5 or later.<br>To determine the current release of the Virtual I/O Server and update it if necessary, see the following instructions:<br>• ioslevel Command in the *Virtual I/O Server and Integrated Virtualization Manager Command Reference*[1]. | X | X |

**Note:**

1. To view the PDF file of the *Virtual I/O Server and Integrated Virtualization Manager Command Reference* (SA76-0101), approximately 4 MB in size, see sa76-0101.pdf .

## Operating system requirements

The following table shows the supported operating system requirements for Partition Mobility.

Table 55. Operating system requirements

| Operating system requirement | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The operating system running in the mobile partition must be AIX or Linux. | X | X |
| The operating system must be at one of the following levels:<br>• AIX 5L Version 5.3 with the 5300-07 Technology Level or later<br>• SUSE Linux Enterprise Server 10 Service Pack 1 or later<br>Earlier versions of AIX and Linux can support an inactive Partition Mobility if the operating systems support virtual devices and POWER6 models. | X | |

## Storage requirements

The following table shows the storage requirements for the source and destination server.

Table 56. Source and destination server storage requirements

| Storage requirements | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The mobile partition must be using storage that is visible to the Virtual I/O Servers on both the source and destination systems. | X | X |

## Preparing for an HMC migration

Use this information to help gain an understanding of what to consider when planning to migrate an active or inactive logical partition using the Hardware Management Console.

### Preparing the source and destination servers for Partition Mobility

You must complete several tasks to prepare the source and destination server for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

To prepare the source and destination server for Partition Mobility, complete the following tasks.

*Table 57. Planning tasks for the source and destination servers*

| Server planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Ensure that the source and destination servers meet the requirements for Partition Mobility. See Requirements for partition mobility using the HMC for information. | X | X |
| 2. Ensure that the source and destination servers are managed by the same HMC (or redundant HMC pair).<br>**Note:** The Validate function in the Partition Mobility wizard checks this for you. | X | X |
| 3. Ensure that the logical memory block size is the same on the source and destination servers. To determine the logical memory block size of each server, and update the sizes if necessary, see Changing the logical memory block size for instructions.<br>**Note:** The Validate function in the Partition Mobility wizard checks this for you. | X | X |
| 4. Ensure that the destination server is not running on battery power. If the destination server is running on battery power, return the server to its regular power source before moving a logical partition. | X | X |
| 5. Ensure that the destination server has enough available memory to support the mobile partition. See "Determining available memory on the destination server" on page 224 for instructions.<br>**Note:** The Validate function in the Partition Mobility wizard checks this for you. | X | |

*Table 57. Planning tasks for the source and destination servers (continued)*

| Server planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 6. Ensure that the destination server has enough available processors to support the mobile partition. See "Determining available processors on the destination server" on page 225 for instructions.<br>**Note:** The Validate function in the Partition Mobility wizard checks this for you. | X | |
| 7. Verify that the source and destination mover service partition can communicate with each other.<br>**Note:** The Validate function in the Partition Mobility wizard checks this for you. | X | |

**Changing the logical memory block size:**

You might enhance the managed system performance by manually or automatically changing the logical memory block size.

The system kernel uses the memory block size to read and write files. By default, the logical memory block size is set to **Automatic**. This setting allows the system to set the logical block memory size based on the physical memory available. You can also manually change the logical memory block size.

To select a reasonable logical block size for your system, consider both the performance desired and the physical memory size. Use the following guidelines when selecting logical block sizes:

- On systems with a small amount of memory installed (2 GB or less), a large logical memory block size results in the firmware consuming an excessive amount of memory. Firmware must consume at least 1 logical memory block. As a general rule, select the logical memory block size to be no greater than 1/8th the size of the system's physical memory.
- On systems with a large amount of memory installed, small logical memory block sizes result in a large number of logical memory blocks. Because each logical memory block must be managed during boot, a large number of logical memory blocks can cause boot performance problems. As a general rule, limit the number of logical memory blocks to 8 K or less.

**Note:** The logical memory block size can be changed at run time, but the change does not take effect until the system is restarted.

To perform this operation, your authority level must be one of the following:

- Administrator
- Authorized service provider

To configure logical memory block size, perform the following task:

1. On the ASMI Welcome pane, specify your user ID and password, and click **Log In**.
2. In the navigation area, expand **Performance Setup**.
3. Select **Logical Memory Block Size**.
4. In the right pane, select the logical memory block size and click **Save settings**.

**Determining available memory on the destination server:**

This procedure provides instructions that explain how to determine the available memory on the destination server and allocate more memory if necessary.

You must be a super administrator to perform this task.

To determine the available memory on the destination server using the HMC, complete the following steps:

1. Determine how much memory the mobile partition requires:
   a. In the navigation area, open **Systems Management** and select **Servers**.
   b. Select the managed server of your choice in the navigation area.
   c. In the contents area, select the logical partition of your choice.
   d. Select **Properties** and select the **Hardware** tab and the **Memory** tab.
   e. View the Memory section and record the minimum, maximum, and available memory settings.
   f. Click **OK**.
2. Determine the memory available on the destination server:
   a. In the navigation area, select **Systems Management** and select **Servers**.
   b. Select the managed server of your choice in the navigation area.
   c. Select **Properties** and the **Memory** tab.
   d. Record the **Current memory available for logical partition usage (MB)** .
   e. Click **OK**.
3. Compare the values from steps 1 and 2.
   - If the destination server has enough available memory to support the mobile partition, continue to "Preparing the source and destination servers for Partition Mobility" on page 223.
   - If the destination server does not have enough available memory to support the mobile partition, use the HMC to dynamically remove memory from the logical partition or you can remove memory from logical partitions on the destination server.

**Determining available processors on the destination server:**

This procedure provides instructions that explain how to determine the available processors on the destination server and allocate more processors if necessary.

You must be a super administrator to perform this task.

To determine the available processors on the destination server using the HMC, complete the following steps:

1. Determine how many processors the mobile partition requires:
   a. In the navigation area, open **Systems Management** and select **Servers**.
   b. Select the managed server of your choice in the navigation area.
   c. In the contents area, select the logical partition of your choice
   d. Select **Properties** and select the **Hardware** tab and the **Processors** tab.
   e. View the Processor section and record the minimum, maximum, and available processor settings.
   f. Click **OK**.
2. Determine the processors available on the destination server:
   a. In the navigation area, open **Systems Management** and select **Servers**.
   b. Select the managed server of your choice in the navigation area.
   c. Select **Properties** and the **Processors** tab.
   d. Record the **Available processors**.
   e. Click **OK**.

3. Compare the values from steps 1 and 2.
   - If the destination server has enough available processors to support the mobile partition, continue to“Preparing the source and destination servers for Partition Mobility” on page 223.
   - If the destination server does not have enough available processors to support the mobile partition, use the HMC, to dynamically remove the processors from the logical partition or you can remove processors from logical partitions on the destination server.

## Preparing the HMC for Partition Mobility

Complete the task to prepare the HMC for Partition Mobility.

To prepare the HMC for Partition Mobility, complete the following task.

*Table 58. Planning task for the HMC*

| HMC planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Ensure that the HMC meets the requirements for Partition Mobility. See “Requirements for Partition Mobility using the Hardware Management Console” on page 219 for information. | X | X |

## Preparing the source and destination Virtual I/O Server logical partitions for Partition Mobility

You must complete several tasks to prepare the source and destination Virtual I/O Server logical partitions for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

*Table 59. Planning tasks for the source and destination Virtual I/O Server logical partitions*

| Virtual I/O Server planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Ensure that the source and destination servers meet the requirements for Partition Mobility. See “Requirements for Partition Mobility using the Hardware Management Console” on page 219 for information. | X | X |
| 2. Ensure that the Mover Service Partition is enabled on one or more source and destination Virtual I/O Server logical partitions. See “Enabling the mover service partitions” for instructions. | X | |
| 3. Optional: Synchronize the time-of-day clocks for the source and destination Virtual I/O Server logical partitions. See “Synchronizing the time-of-day clocks” on page 227 for instructions. | X | |

**Enabling the mover service partitions:**

This procedure provides instructions that explain how to enable the mover service partition using the Hardware Management Console (HMC).

You must be a super administrator or operator to complete this task.

There must be at least one mover service partition on the source and destination Virtual I/O Server for the mobile partition to participate in active Partition Mobility. If the mover service partition is disabled on either the source or destination Virtual I/O Server, the mobile partition can participate only in inactive Partition Mobility.

To enable the source and destination mover service partition using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. In the contents area, select a Virtual I/O Server logical partition and select **Properties**.
4. On the **General** tab, select **Mover Service Partition**, and click **OK**.
5. Repeat steps 3 and 4 for the destination server.

**Synchronizing the time-of-day clocks:**

This optional procedure provides instructions that explain how to synchronize the time-of-day clocks for the source and destination Virtual I/O Server logical partitions.

You must be a super administrator to complete this task.

Synchronizing the time-of-day clocks for the source and destination Virtual I/O Server logical partitions is an optional step for active partition mobility. If you choose not to complete this step, the source and destination servers will synchronize the clocks while the mobile partition is moving from the source server to the destination server. Completing this step before the mobile partition is moved can prevent possible errors.

To synchronize the time-of-day clocks on the source and destination Virtual I/O Servers logical partition using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. In the contents area, select a Virtual I/O Server logical partition and select **Properties**.
4. Click the **Settings** tab.
5. Select **Enable** for Time reference and click **OK**.
6. Repeat steps 3 through 5 for the destination server and the destination Virtual I/O Server.

## Preparing the mobile partition for Partition Mobility

You must complete several tasks to prepare the mobile partition for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

To prepare the mobile partition for Partition Mobility, complete the following tasks.

*Table 60. Planning tasks for the mobile partition*

| Mobile partition planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Ensure that the operating system meets the requirements for Partition Mobility. See “Requirements for Partition Mobility using the Hardware Management Console” on page 219 for information.<br>**Note:** The mobile partition cannot be a Virtual I/O Server logical partition. | X | X |

*Table 60. Planning tasks for the mobile partition (continued)*

| Mobile partition planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 2. Ensure that Resource Monitoring and Control (RMC) connections are established with the mobile partition, the source and destination Virtual I/O Servers, and the source and destination mover service partitions. See "Verifying Resource Monitoring and Control connections for the mobile partition" on page 229 for instructions. | X | |
| 3. Ensure that the mobile partition is not enabled for redundant error path reporting. See "Disabling the mobile partition for redundant error path reporting" on page 230 for instructions. | X | X |
| 4. Ensure that the mobile partition is only using a virtual serial adapter for virtual terminal connections. See "Disabling virtual serial adapters for the mobile partition" on page 230 for instructions. | X | X |
| 5. Ensure that the mobile partition is not part of a partition workload group. See "Removing the mobile partition from a partition workload group" on page 231 for instructions. | X | X |
| 6. Ensure that the mobile partition is not using barrier synchronization register (BSR) arrays. See "Disabling BSR for the mobile partition" on page 231 for instructions. | X | |
| 7. Ensure that the mobile partition is not using huge pages. See "Disabling huge pages for the mobile partition" on page 232 for instructions. | X | |
| 8. Ensure that the mobile partition does not have physical I/O adapters or aHost Ethernet Adapter (or Integrated Virtual Ethernet). See "Removing dedicated I/O from the mobile partition" on page 233 and "Removing Host Ethernet Adapters from the mobile partition" on page 233 for instructions. | X | |
| 9. (Optional) Determine the name of the partition profile for the mobile partition on the destination server. | X | X |
| 10. Ensure that the applications running in the mobile partition are mobility-safe or mobility-aware. See Software applications that recognize migrations for more information. | X | |

*Table 60. Planning tasks for the mobile partition (continued)*

| Mobile partition planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 11. If you changed any partition profile attributes, shut down and activate the new profile for the new values to take effect. Use the following steps to complete this task.<br>1. Shut down the mobile partition. See Shutting down an operating system in the *Operations Guide for the Hardware Management Console and Managed Systems*[1].<br>2. Activate the partition profile of the mobile partition. See Activating a partition profile for instructions. | X | X |
| **Note:**<br>1. To view the PDF file of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), approximately 4 MB in size, see sa76-0085.pdf. | | |

**Verifying Resource Monitoring and Control connections for the mobile partition:**

This procedure provides instructions that explain how to verify a Resource Monitoring and Control (RMC) connection for the mobile partition.

You must be a super administrator to complete this task.

RMC is a no-charge feature of AIX that can be configured to monitor resources and perform an action in response to a defined condition. With RMC, you can configure response actions or scripts that manage general system conditions with little or no involvement from the system administrator. On the HMC, RMC is being used as the main communication channel between AIX and Linux logical partitions and the HMC.

To verify an RMC connection for the mobile partition, complete the following steps:

1. Using the HMC command line, enter `lspartition -dlpar` .

   Your command results will look similar to this example:

   - If the results for your logical partition are `<Active 1>`, then the RMC connection is established. Skip the rest of this procedure and return to “Preparing the mobile partition for Partition Mobility” on page 227.
   - If the results for your logical partition are `<Active 0>` or your logical partition is not displayed in the command results, continue to the next step.
2. Verify that the RMC firewall port on the HMC is disabled.
   - If the RMC firewall port is disabled, skip to step 3.
   - If the RMC firewall port is enabled, change your HMC firewall setting. Repeat step 1.
3. Use telnet to access the logical partition. If you cannot use telnet, open a virtual terminal on the HMC to set up the network on the logical partition.
4. If the logical partition network has been set up correctly and there is still no RMC connection, verify that the RSCT fileset is installed.

- If the RSCT fileset is installed, use telnet to the HMC from the logical partition to verify if the network is working correctly and that the firewall has been disabled. After verifying these tasks, repeat step 1. If you continue to have problems establishing an RMC connection for your mobile partition, contact your next level of support.
- If the RSCT fileset is not installed, use your AIX installation CD to install the fileset.

**Note:** It takes approximately five minutes for RMC connection to establish the connection after the network setup has been changed or after activating the logical partition.

**Disabling the mobile partition for redundant error path reporting:**

This procedure provides instructions that explain how to disable the mobile partition for redundant error path reporting.

You must be a super administrator to complete this task.

Redundant error path reporting allows a logical partition to report common server hardware problems and logical partition hardware errors to the HMC. If you want to migrate a logical partition, disable the redundant error path reporting.

To disable the mobile partition for redundant error path reporting using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. In the contents area, select the logical partition of your choice.
4. Select **Configuration > Manage Profiles.**
5. Select the profile of your choice and select **Actions > Edit**.
6. Click the **Settings** tab.
7. Deselect **Enable redundant error path reporting** and click **OK**. For this change to take effect, activate this logical partition with this profile.

**Disabling virtual serial adapters for the mobile partition:**

This procedure provides instructions that explain how to disable unreserved virtual serial adapters for the mobile partition.

You must be a super administrator to complete this task.

Virtual serial adapters are often used for virtual terminal connections to the operating system. The first two virtual serial adapters (slots 0 and 1) are reserved for the HMC. For a logical partition to participate in Partition Mobility , it cannot have any virtual serial adapters, except for the two that are reserved for the HMC.

To disable unreserved virtual serial adapters using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. In the contents area, select the logical partition of your choice.
4. Select **Configuration > Manage Profiles.**
5. Select the profile of your choice and select **Actions > Edit**.
6. Select the **Virtual Adapter** tab.
7. If there are more than two virtual serial adapters listed, then ensure that the additional adapters beyond 0 and 1 are not selected as **Required**.

- If you have additional virtual serial adapters listed as **Required**, select the adapter that you would like to remove. Then select **Actions > Delete** to remove the adapter from the partition profile.
- You can select **Dynamic Logical Partitioning > Virtual Adapters**. The Virtual Adapters panel is displayed. Select the adapter that you would like to remove and select **Actions > Delete** to remove the adapter from the partition profile.

8. Click **OK**.

**Removing the mobile partition from a partition workload group:**

This procedure provides instructions that explain how to remove the mobile partition from a partition workload group.

You must be a super administrator to complete this task.

A partition workload group identifies a set of logical partitions that are located on the same physical system. The partition profile specifies the name of the partition workload group that it belongs to, if applicable. A partition workload group is defined when you use the HMC to configure a logical partition. For a logical partition to participate in Partition Mobility, it cannot be assigned to a partition workload group.

To remove the mobile partition from a partition workload group using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. In the contents area, select the logical partition of your choice.
4. Select **Configuration > Manage Profiles.**
5. Select the profile of your choice and select **Actions > Edit**.
6. Click the **Settings** tab.
7. In the Workload Management area, select **(None)** and click **OK**.
8. Repeat steps 1 through 7 for all partition profiles associated with the mobile partition. For this change to take effect, you will need to activate this logical partition with this profile.

**Disabling BSR for the mobile partition:**

This procedure provides instructions that explain how to disable barrier synchronization register (BSR) arrays for the mobile partition.

You must be a super administrator to perform this task.

BSR is a memory register that is located on certain POWER processor-based systems. A parallel-processing application running on AIX can use a BSR to perform barrier synchronization, which is a method for synchronizing the threads in the parallel-processing application.

For a logical partition to participate in active Partition Mobility, it cannot use BSR arrays. If the mobile partition uses BSR, the logical partition can participate in inactive Partition Mobility.

To disable BSR for the mobile partition using the HMC, complete the following steps:

1. In the navigation area, select **Systems Management** and select **Servers**.
2. In the navigation area, select the managed server of your choice and select **Properties**.
3. Click the **Capabilities** tab.
   - If barrier synchronization register (BSR) Capable is **True**, click **OK** and continue with the next step.

- If barrier synchronization register (BSR) Capable is **False**, the server does not support BSR. Skip the rest of this procedure and continue to "Preparing the mobile partition for Partition Mobility" on page 227.

4. In the navigation area, open **Systems Management** and select **Servers**.
5. Select the managed server of your choice in the navigation area.
6. Select the logical partition of your choice in the contents area.
7. Select **Properties** in the Task area.
8. Click the **Hardware** tab.
9. Click the **Memory** tab.
   - If the number of BSR arrays equals zero, the mobile partition can participate in active or inactive Partition Mobility. Skip the rest of this procedure and continue to "Preparing the mobile partition for Partition Mobility" on page 227.
   - If the number of BSR arrays is not equal to zero, then take one of the following actions:
     - Perform an inactive movement instead of an active movement.
     - Click **OK** and continue to the next step to prepare the mobile partition for an active movement.
10. Select the mobile partition, and then select **Configuration > Manage Profiles**.
11. Select the partition profile with which you will reactivate the mobile partition, and select **Action > Edit**.
12. Click the **Memory** tab.
    - If the number of BSR arrays equals 0, the mobile partition can participate in active or inactive Partition Mobility. Skip the rest of this procedure and continue to "Preparing the mobile partition for Partition Mobility" on page 227.
    - If the number of BSR arrays is not equal to 0, then take the following action to change BSR to 0 if you want to do an active migration:
      - Enter 0 in the field for the BSR arrays.
      - Click **OK** and continue to the next step to prepare the mobile partition for an active movement.
13. Activate this logical partition with this profile in order for this change to take effect.

**Disabling huge pages for the mobile partition:**

This procedure provides instructions that explain how to disable huge pages for the mobile partition.

You must be a super administrator to perform this task.

Huge pages can improve performance in specific environments that require a high degree of parallelism, such as in DB2 partitioned database environments. You can specify the minimum, desired, and maximum number of huge pages to assign to a logical partition when you create the logical partition or partition profile.

For a logical partition to participate in active Partition Mobility, it cannot use huge pages. If the mobile partition uses huge pages, it can participate in inactive Partition Mobility.

To disable huge pages for the mobile partition using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select **Properties** in the Task area.
4. Click the **Capabilities** tab.
   - If Huge Page Capable is **True**, then click **OK** and continue with the next step.

- If Huge Page Capable is **False**, then the source server does not support huge pages. The mobile partition can participate in active or inactive Partition Mobility. Skip the rest of this procedure and continue to "Preparing the mobile partition for Partition Mobility" on page 227.

5. In the navigation area, open **Systems Management** and select **Servers**.
6. Select the managed server of your choice in the navigation area.
7. In the contents area, select the logical partition of your choice.
8. Select **Properties** and the **Hardware** tab and then the **Memory** tab.
   - If the current huge page memory equals 0, then skip the rest of this procedure and continue to "Preparing the mobile partition for Partition Mobility" on page 227.
   - If the current huge page memory is not equal to 0, then take one of the following actions:
     - – Perform an inactive movement instead of an active movement.
     - – Click **OK** and continue with the next step to prepare the mobile partition for an active movement.
9. In the navigation area, open **Systems Management** and select **Servers**.
10. Select the managed server of your choice in the navigation area.
11. In the contents area, select the logical partition of your choice.
12. Select **Configuration > Manage Profiles.**
13. Select the profile of your choice and select **Actions > Edit**.
14. Click the **Memory** tab.
15. Enter **0** in the field for desired huge page memory, and click **OK**.
16. Activate this logical partition with this profile in order for this change to take effect.

**Removing dedicated I/O from the mobile partition:**

This procedure provides instructions that explain how to remove dedicated I/O from the mobile partition.

You must be a super administrator to perform this task.

Dedicated I/O devices are physical devices found in the server unit itself and in expansion units and towers that are attached to the server. For a logical partition to participate in active Partition Mobility, it cannot have dedicated I/O. All I/O must be virtual. You can use inactive Partition Mobility if a mobile partition has dedicated I/O. The I/O devices will automatically be removed from the logical partition before the migration occurs.

To remove dedicated I/O from the mobile partition using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select the logical partition of your choice in the contents area.
4. Select **Configuration > Manage Profiles** .
5. Select the partition profile of your choice and select **Actions > Edit**.
6. Select the **I/O** tab.
   - If **Required or Desired** is selected for any resource, take one of the following actions:
     - – Select the required or desired devices and click **Remove** and click **OK**. Continue with the next step to prepare the mobile partition for an active partition migration.
7. Shut down the mobile partition, then power it on using the profile with the dedicated I/O resource modifications you just made.

**Removing Host Ethernet Adapters from the mobile partition:**

This procedure provides instructions that explain how to remove a Host Ethernet Adapter (or Integrated Virtual Ethernet), from a mobile partition using the HMC.

You must be a super administrator to perform this task.

To remove a Host Ethernet Adapter from the mobile partition using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select the mobile partition and select **Configuration > Manage Profiles**.
4. Select the partition profile of your choice and select **Actions > Edit**.
5. Select the **Logical Host Ethernet Adapters (LHEA)** tab.
6. Select the physical port locations that have a logical port ID assigned to it and click **Reset**.
7. Click **OK**.

**Activating a logical partition using the HMC:**

You must activate a logical partition before you can use the logical partition. When you activate a logical partition, the system commits resources to the logical partition and starts the operating system or software that is installed on the logical partition.

When you activate a logical partition, you must select a partition profile. A *partition profile* is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition.

To activate a logical partition using the HMC, you must be a super administrator, operator, or product engineer. For more information about user roles, refer to Tasks and roles in the *Operations Guide for the Hardware Management Console and Managed Systems*. To view the abstract of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), see sa76-0085.pdf

To activate a logical partition using the HMC, follow these steps:

1. In the navigation pane, open **Systems Management**, open **Servers**, and click the system on which the logical partition is located.
2. In the contents pane, select the logical partition, click the **Tasks** button, and choose **Operations** → **Activate**.
3. Select the partition profile that you want to use to activate the logical partition.
4. If you want the HMC to open a terminal window or console session for the logical partition when the logical partition is activated, select **Open a terminal window or console session**.
5. If you want to use a keylock position or boot mode that is different from the keylock position or boot mode specified in the partition profile, click **Advanced**, select the desired keylock position and boot mode, and click **OK**.
6. Click **OK**.

## Preparing the storage configuration for Partition Mobility

You must complete several tasks to ensure your storage configuration meets the minimal configuration for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

If your environment already meets the minimal configuration, then select verify Verifying the existing storage configuration.

If you need to set up this environment, then select Configuring the storage environment.

### Verifying the existing storage configuration for Partition Mobility:

Complete the tasks to verify that your storage configuration meets the minimal configuration requirements for Partition Mobility.

Before you migrate your logical partition, verify that your existing storage configuration meets the requirements needed to migrate your logical partition. Use the following tasks to understand how your storage configuration is set up.

Verify the following attributes, assignments, and connections.

*Table 61. Planning tasks for storage*

| Storage planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Verify that the reserve_policy attributes on the physical volumes are set to no_reserve. | X | X |
| 2. Verify that the virtual devices have the same unique identifier, physical identifier, or an IEEE volume attribute. See Identifying exportable disks | X | X |
| 3. Verify that the mobile partition has access to the source Virtual I/O Server virtual SCSI adapter. See "Verifying that the mobile partition will have access to the destination Virtual I/O Server virtual SCSI adapter" on page 238 for instructions. | X | X |
| 4. Verify that the mobile partition will have access to the destination Virtual I/O Server virtual SCSI adapter after it moves to the destination system. See "Verifying that the mobile partition will have access to the destination Virtual I/O Server virtual SCSI adapter" on page 238 for instructions. | X | X |
| 5. Verify that the mobile partition has access to the physical storage. See "Verifying that the mobile partition has access to its physical storage" on page 244 for instructions. | X | X |
| 6. Verify that the mobile partition does not have physical or dedicated I/O adapters and devices. See "Removing dedicated I/O from the mobile partition" on page 233 for instructions. | X | |

*Table 61. Planning tasks for storage (continued)*

| Storage planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 7. If you changed any partition profile attributes, complete the following steps in order for the new values to take effect.<br>1. Shut down the logical partition associated with the changed logical partition profile. See Shutting down an operating system in the *Operations Guide for the Hardware Management Console and Managed Systems* for instructions[1].<br>2. Activate the changed partition profile. See Activating a partition profile for instructions. | X | X |
| **Note:**<br>1. To view the PDF file of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), approximately 4 MB in size, see sa76-0085.pdf . | | |

*Identifying exportable disks:*

To export a physical volume as a virtual device, the physical volume must have an IEEE volume attribute, a unique identifier (UDID), or a physical identifier (PVID).

To identify exportable disks, complete the following steps:

1. Determine whether a device has an IEEE volume attribute identifier by running the following command from the Virtual I/O Server command line:

   ```
   lsdev -dev hdiskX -attr
   ```

   Disks with an IEEE volume attribute identifier have a value in the `ieee_volname` field. Output similar to the following is displayed:

   ```
   ...
   cache_method   fast_write                       Write Caching method
      False
   ieee_volname   600A0B800012DD0D00000AB441ED6AC IEEE Unique volume name
      False
   lun_id         0x001a000000000000               Logical Unit Number
      False
   ...
   ```

   If the `ieee_volname` field does not appear, then the device does not have an IEEE volume attribute identifier.

2. If the device does not have an IEEE volume attribute identifier, then determine whether the device has a UDID by completing the following steps:
   a. Type `oem_setup_env`.
   b. Type `odmget -qattribute=unique_id CuAt`. The disks that have a UDID are listed. Output similar to the following is displayed:

      ```
      CuAt:
       name = "hdisk1"
       attribute = "unique_id"
       value = "2708ECVBZ1SC10IC35L146UCDY10-003IBMscsi"
       type = "R"
      ```

```
 generic = ""
 rep = "nl"
 nls_index = 79

CuAt:
 name = "hdisk2"
 attribute = "unique_id"
 value = "210800038FB50AST373453LC03IBMscsi"
 type = "R"
 generic = ""
 rep = "nl"
 nls_index = 79
```

Devices in the list that are accessible from other Virtual I/O Server partitions can be used in virtual SCSI MPIO configurations.

c. Type exit.

3. If the device does not have either an IEEE volume attribute identifier or a UDID, then determine whether the device has a PVID by running the following command:

```
lspv
```

The disks and their respective PVIDs are listed. Output similar to the following is displayed:

```
NAME        PVID                   VG       STATUS
hdisk0      00c5e10c1608fd80       rootvg   active
hdisk1      00c5e10cf7eb2195       rootvg   active
hdisk2      00c5e10c44df5673       None
hdisk3      00c5e10cf3ba6a9a       None
hdisk4      none                   None
```

4. If the device does not have either an IEEE volume attribute identifier, a UDID, or a PVID, then complete one of the following tasks to assign an identifier:
   a. Upgrade your vendor software and then repeat this entire procedure, Identifying exportable disks, from the beginning. The latest versions of some vendor software include support for identifying devices using a UDID. Before upgrading, ensure that you preserve any virtual SCSI devices that you created when using the versions of the software that did not support identifying devices using a UDID. For information and upgrade instructions, see the documentation provided by your vendor software.
   b. If the upgraded vendor software does not produce a UDID or IEEE volume attribute identifier, then put a PVID on the physical volume by running the following command:

```
chdev -dev hdiskX -attr pv=yes
```

*Changing the source Virtual I/O Server virtual SCSI adapter properties:*

This procedure provides instructions that explains how to change the connection preferences and slot assignment of the source Virtual I/O Server virtual SCSI adapter.

If you are unsure about the slot assignments and connection preferences for the source Virtual I/O Server virtual SCSI adapter, complete the "Virtual I/O adapter worksheet for partition mobility" on page 242

You must be a super administrator to perform this task.

To change the properties of the virtual SCSI adapter using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select the Virtual I/O Server of your choice in the contents area.
4. Select **Configuration > Manage Profiles**.
5. Select the partition profile of your choice and select **Actions > Edit**.
6. Select the **Virtual Adapters** tab.

7. Select the **Server SCSI** device and click **Actions > Edit**.
   a. In the **Adapter** field, enter the slot number of the virtual SCSI adapter created in the Virtual I/O Server for the mobile partition.
   b. Select **Only selected client logical partition can connect**.
   c. In the **Client partition** field, select the client logical partition.
   d. In the **Client adapter ID** field, enter the client adapter ID associated with the client logical partition.
   e. Click **OK** to save the changes to the partition profile.
8. Click **OK** to exit the Partition Profile Properties window.

*Changing the virtual SCSI adapter properties of the mobile partition:*

This procedure provides instructions that explains how to change the connection preferences and slot assignments of the source virtual SCSI adapter on the mobile partition.

If you are unsure about the slot assignments and connection preferences for the virtual SCSI adapters on the mobile partition, complete the "Virtual I/O adapter worksheet for partition mobility" on page 242.

To change the properties of a virtual SCSI adapter, you must be a super administrator.

To change the properties of the virtual SCSI adapter using the HMC, complete the following steps:
1. In the navigation area, open **Systems Management**.
2. Select **Server**.
3. Select a mobile partition and select **Properties**.
4. Select the **Virtual Adapters** tab.
5. Expand **Virtual SCSI** and select the virtual SCSI adapter that you want to access the virtual SCSI adapter on the source Virtual I/O Server logical partition.
6. Open the source Virtual I/O Server.
7. Select the slot number of the Virtual I/O Server virtual SCSI adapter to which you want the mobile partition to access and click **OK**.

*Verifying that the mobile partition will have access to the destination Virtual I/O Server virtual SCSI adapter:*

This procedure provides step-by-step instructions that explains how to verify that the mobile partition will have access to the Virtual I/O Server virtual SCSI adapter after the mobile partition is moved to the destination system.

To verify the virtual SCSI configuration, you must be a super administrator.

The virtual SCSI adapter on the Virtual I/O Server provides client logical partitions (including the mobile partition) access to storage. For partition mobility (active or inactive) to be successful the source Virtual I/O Server logical partitions must have at least one virtual SCSI adapter configured to allow access to the mobile partition.

To verify the virtual SCSI configuration using the HMC, complete the following steps:
1. Verify the virtual SCSI adapter configuration of the destination Virtual I/O Server virtual SCSI adapter:
   a. In the navigation area, open **Systems Management** and select **Servers**.
   b. Select the destination server of your choice in the navigation area.
   c. Select the logical partition of your choice in the contents area.
   d. Select **Properties** in the Task area.

e. Select **Virtual Adapters** tab.
f. Expand **Virtual SCSI** .
g. Verify that the **Remote partition** and the **Remote Adapter** are both blank.

2. If the values are not blank, then complete the following steps:
   a. In the navigation area, open **Systems Management** and select **Servers**.
   b. Select the managed server of your choice in the navigation area.
   c. Select the logical partition of your choice in the contents area.
   d. Select **Properties** in the Task area.
   e. Select the **Virtual Adapters** tab.
   f. Click **Actions > Create > SCSI Adapter**. The SCSI Adapter create window is shown.
   g. Select **Any client logical partition can connect** and click **OK**.
   h. Click **Submit** to exit the **SCSI Client Adapter Properties** window.
   i. Click **OK** to save the changes to the logical partition profile.

**Configuring the storage environment for Partition Mobility:**

You must complete several tasks to configure your storage environment for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

To set up your storage configuration, complete the following tasks.

*Table 62. Planning tasks for storage*

| Storage planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Set up and configure a storage area network. See Virtual I/O Server Supported Environment for a list of storage devices supported by the Virtual I/O Server. | X | X |
| 2. Configure a SCSI adapter on the source Virtual I/O Servers. See "Adding a SAN Controller to a Virtual I/O Server" on page 240 for instructions. | X | X |
| 3. Connect the source and destination Virtual I/O Server logical partitions to the storage area network. | X | X |
| 4. Set the reserve_policy attributes on the mobile client physical volumes to no_reserve | X | X |
| 5. Assign the same disks to the source and all possible destination Virtual I/O Server logical partitions. | X | X |
| 6. Determine the virtual SCSI adapter slot assignments and connection specifications. See "Virtual I/O adapter worksheet for partition mobility" on page 242 for instructions. | X | X |

*Table 62. Planning tasks for storage (continued)*

| Storage planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 8. Create a connection between the SCSI adapter and the virtual SCSI adapter on the source Virtual I/O Server. See Creating the virtual target device on the Virtual I/O Server for instructions. | X | X |
| 9. Configure the virtual SCSI adapters on the mobile partition. See "Configuring a virtual SCSI adapter on the mobile partition" on page 243 for instructions. | X | X |
| 10. Activate the mobile partition to establish communication between the virtual SCSI adapter on mobile partition and the virtual adapter on the source Virtual I/O Server logical partition. See Activating a partition profile for instructions. | X | X |
| 11. Verify that the mobile partition has access to the physical storage. See "Verifying that the mobile partition has access to its physical storage" on page 244 for instructions. | X | X |
| 12. Verify that the mobile partition does not have physical or dedicated I/O adapters and devices. See "Removing dedicated I/O from the mobile partition" on page 233 for instructions. | X | |
| 13. If you changed any partition profile attributes, complete the following steps in order for the new values to take effect.<br>1. Shut down the logical partition associated with the changed partition profile. See Shutting down an operating system in the *Operations Guide for the Hardware Management Console and Managed Systems* for instructions[1].<br>2. Activate the changed partition profile. See Activating a partition profile for instructions. | X | X |
| **Note:**<br>1. To view the PDF file of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), approximately 4 MB in size, see sa76-0085.pdf . | | |

*Adding a SAN Controller to a Virtual I/O Server:*

This procedure provides instructions that explain how to add at least one SAN Controller on the source and destination Virtual I/O Servers.

You must be a super administrator to complete this task.

The SAN Controller on the Virtual I/O Server provides the mobile partition with access to storage. For Partition Mobility (active or inactive) to be successful, both the source and destination Virtual I/O Server logical partitions must have a SAN Controller and access to the same mobile partition storage.

To add a SAN Controller using the HMC, complete the following steps:

1. Dynamically add a physical Fibre Channel adapter to the source Virtual I/O Server logical partition. See Dynamically adding physical I/O devices and slots for instructions:
2. Update the partition profiles of the source Virtual I/O Server with the new SAN Controller. You must update the partition profiles so that the next time they are activated, the new SAN Controller is not lost.
   a. In the navigation area, open **Systems Management** and select **Servers**.
   b. Select the managed server of your choice in the navigation area.
   c. Select the Virtual I/O Server of your choice in the contents area.
   d. Select **Configuration > Manage Profiles**.
   e. Select the partition profile of your choice and select **Actions > Edit**.
   f. Select the **I/O** tab.
   g. Select a Fibre Channel adapter from the **Physical I/O** table and click **Add as required**. The device appears in the **Added** column and is marked as **Required**.
   h. Click **OK** to save the changes to the partition profile.
3. Repeat this procedure for the destination server and the destination Virtual I/O Server logical partition.

*Adding physical I/O devices and slots dynamically using the HMC:*

You can add a physical I/O slot (and the adapter and devices that are connected to that slot) to a running logical partition using the Hardware Management Console (HMC). This allows you to add I/O capabilities to a running logical partition without having to shut down the logical partition.

A Linux logical partition supports the dynamic addition of physical I/O slots only if the following conditions are met:

- A Linux distribution that supports dynamic logical partitioning is installed on the Linux logical partition. Distributions that support dynamic logical partitioning include SUSE Linux Enterprise Server 9 and later versions.
- The DynamicRM tool package is installed on the Linux logical partition. For more information on the DynamicRM tool package, see the Service and productivity tools Web site.

To add a physical I/O slot dynamically to a running logical partition using the HMC, you must be a super administrator, service representative, product engineer, or operator. For more information about user roles, refer to Tasks and roles in the *Operations Guide for the Hardware Management Console and Managed Systems*. To view the abstract of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), see sa76-0085.pdf .

To add a physical I/O slot dynamically to a running logical partition using the HMC, follow these steps:

1. In the navigation pane of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition is located.
2. In the contents pane, select the logical partition, click the **Tasks** button, and choose **Dynamic Logical Partitioning** → **Physical Adapters** → **Add**.
3. Select the physical I/O slot that you want to add to the logical partition.

4. If you want to assign the physical I/O slot to an I/O pool, select the I/O pool for the physical I/O slot in **I/O Pool ID**.
5. Click **OK**.

*Virtual I/O adapter worksheet for partition mobility:*

Use this worksheet to help you plan the slot assignments and connection specifications of the virtual I/O adapters in your environment.

For the mobile partition to access storage, it must have access to a Virtual I/O Server virtual SCSI adapter. You can specify that only the mobile partition have access to a particular virtual SCSI server adapter. This specification requires knowledge of the slot assignment of the virtual SCSI client adapter on the mobile partition.

Use the following table to help you plan the slot assignments and connection specifications for the source and destination Virtual I/O Server virtual I/O adapters and the virtual I/O adapter on the mobile partition. To determine the slots available for a particular adapter, see “Determining available virtual SCSI adapter slots.”

*Table 63. Worksheet for the slot assignments and connection specifications of the virtual I/O adapters*

| Virtual I/O adapter | Slot number | Connection specification |
|---|---|---|
| Source Virtual I/O Server virtual SCSI adapter | | |
| Mobile partition virtual I/O adapter on source system | | |

For example, the source Virtual I/O Server has one virtual SCSI adapter in slot 12. The virtual SCSI adapter on the mobile partition was configured in slot 4. The source Virtual I/O Server virtual SCSI adapter was configured to allow access to only the virtual SCSI adapter on the mobile partition. The virtual SCSI adapter on the mobile partition was configured to access the source Virtual I/O Server virtual SCSI adapter.

*Table 64. Example slot assignments and connection specification*

| Virtual I/O adapter | Slot number | Connection specification |
|---|---|---|
| Source Virtual I/O Server virtual SCSI adapter | 12 | Slot 4, which is the virtual SCSI adapter on the mobile partition |
| Mobile partition virtual I/O adapter on source system | 4 | Slot 12, which is the source Virtual I/O Server virtual SCSI adapter |

*Determining available virtual SCSI adapter slots:*

This procedure provides instructions that explain how to determine the available slots for the virtual SCSI adapters on the source Virtual I/O Server, the destination Virtual I/O Server, and the mobile partition.

You must be a super administrator to complete this task.

To view the properties of a virtual SCSI adapter using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select the logical partition of your choice in the contents area.
4. Select **Properties** in the Task area.
5. Select the **Virtual Adapter** tab and click **Virtual SCSI**.

6. View the current slot assignments and click **OK**.

*Configuring a Virtual I/O Server virtual SCSI adapter:*

This procedure provides instructions that explain how to configure at least one virtual SCSI adapter on the source and destination Virtual I/O Servers.

You must be a super administrator to complete this task.

Retrieve your completed “Virtual I/O adapter worksheet for partition mobility” on page 242. You will need this information to configure the source and destination Virtual I/O Server virtual SCSI adapters.

The virtual SCSI adapter on the Virtual I/O Server provides client logical partitions (including the mobile partition) access to storage. For Partition Mobility (active or inactive) to be successful, the source Virtual I/O Server logical partitions must have at least one virtual SCSI adapter configured to allow access to the mobile partition.

To configure a virtual SCSI adapter using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select the source Virtual I/O Server and select **Dynamic Logical Partitioning** → **Virtual Adapters** .
4. Select **Actions > Create > SCSI Adapter**.
   a. In the **Adapter** field, enter the slot number of the virtual SCSI adapter created in the Virtual I/O Server for the mobile partition.
   b. Select **Only selected client partition can connect**.
   c. In the **Client partition** field, select the client logical partition.
   d. In the **Client adapter ID** field, enter the client adapter ID associated with the client logical partition.
   e. Click **OK** to save the changes to the SCSI adapter.
   f. Click **OK** to save the new virtual adapter configuration.
5. Update the partition profiles of the Virtual I/O Server logical partitions with the new virtual I/O adapters. The partition profiles must be updated to ensure that the next time they are activated the new virtual I/O adapters are not lost.
   a. . Select Virtual I/O Server logical partition and select **Configuration -> Manage Profiles**.
   b. Select **Actions > Edit**.
   c. Select the **Virtual I/O Adapters** tab.
   d. Select **Actions > Create > SCSI Adapter**. The SCSI Server Adapter Properties window is shown.
   e. Fill in the fields using the same data from step 4.
   f. Click **OK** to exit the SCSI Server Adapter Properties window.
   g. Click **OK** to save the changes to the partition profile.
   h. Click **OK** to exit the Partition Profile Properties window.

*Configuring a virtual SCSI adapter on the mobile partition:*

This procedure provides instructions that explain how to configure at least one virtual SCSI adapter on the mobile partition.

You must be a super administrator to complete this task.

Retrieve your completed “Virtual I/O adapter worksheet for partition mobility” on page 242. You will need this information to configure the virtual SCSI adapter on the mobile partition.

The virtual SCSI adapter on the mobile partition allows the mobile partition to access storage through its connection to the virtual SCSI adapter on the Virtual I/O Server logical partition. The virtual SCSI adapter on the Virtual I/O Server connects to the physical SCSI adapter on the source Virtual I/O Server which, connects to the storage area network. Ultimately, through these connections, the mobile partition gains access to the physical storage.

To configure a virtual SCSI adapter using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select the mobile partition and select **Dynamic Logical Partitioning > Virtual Adapters**.
4. Select **Actions > Create > SCSI Adapter**. The Create Virtual SCSI Adapter window is shown.
   a. In the **Adapter** field, enter the slot number of the virtual SCSI adapter referenced in the Virtual I/O Server SCSI device for this mobile partition.
   b. In the **Server partition** field, select the Virtual I/O Server.
   c. In the **Server adapter ID** field, enter the server adapter ID associated with the client logical partition
   d. Click **OK** to save the changes to the SCSI adapter.
   e. Click **OK** to save the new virtual adapter configuration.
5. Update the partition profiles of the mobile partition with the new virtual I/O adapters. The partition profiles must be updated. This ensures that the next time they are activated the profiles are activated the new virtual I/O adapters are not lost.
   a. . Select the mobile partition and select **Configuration -> Manage Profiles.**
   b. Select **Actions > Edit**.
   c. Select the **Virtual I/O Adapters** tab.
   d. Select **Actions > Create > SCSI Adapter**. The SCSI Server Adapter Properties window is shown.
   e. Fill in the fields using the same data from step 4.
   f. Click **OK** to exit the **SCSI Client Adapter Properties** window.
   g. Click **OK** to save the changes to the partition profile.
   h. Click **OK** to exit the Partition Profile Properties window.

*Linking virtual SCSI Adapters to the SAN Disk:*

Learn how to link virtual SCSI adapters to the SAN Disk.

1. Using the Virtual I/O Server command line, type `oem_setup_env` then type`lscfg`. Find the hdisks attached to the SAN Controller which will be used for the mobile partition. The hdisks will be listed as MPIO or Disk Array Devices.

   This is an example of a Disk Array hdisk: `+ hdisk3           U787B.001.DNW5D5B-P1-C3-T1-W200900A0B80FD390-L1000000000000  1722-600 (600) Disk Array Device`
2. Type `lsmap -all`. The number after the `C` in the **Physloc** column is the virtual SCSI slot number of the Virtual I/O Server. In this example, the Physloc column details are U9133.55A.100EDCA-V2-C6. *C6* is virtual SCSI slot 6 in this example.
3. Type `mkdev -V hdiskX -vadapter vhostX` where: *X* is the number of the hdisk and vhost you want to link together.
4. Type `exit`, then type `lsmap -all.` A VTD, LUN, Backing device, and Physloc will be listed under the vhost adapter.

*Verifying that the mobile partition has access to its physical storage:*

This procedure provides instructions that explain how to verify that the mobile partition has access to the physical storage on the storage area network.

For Partition Mobility to be successful, the mobile partition must have access to the same physical storage from both the source and destination environments. In the source environment, the following connections must exist:

- Each virtual SCSI adapter on the mobile partition must have access to a target virtual SCSI adapter on the source Virtual I/O Server logical partition.
- The target virtual SCSI adapters on the source Virtual I/O Server logical partition must have access to a SAN host-attached adapter on the source Virtual I/O Server logical partition.
- The SAN host-attached adapter on the source Virtual I/O Server logical partition must be connected to a storage area network and have access to the physical storage devices you want the mobile partition to have access to in the storage area network.

In the destination environment, the following connections must exist:

- The destination Virtual I/O Server logical partition has unused Virtual slots available.
- The SAN host-attached adapter on the destination Virtual I/O Server logical partition must be connected to the same storage area network as the source Virtual I/O Server logical partition and have access to the same mobile partition physical storage as the source Virtual I/O Server logical partition.

You must be a super administrator to complete this task.

To verify these connections using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select the source Virtual I/O Server in the contents area.
4. Select **Hardware (Information) >Virtual Adapters > SCSI** in the Task area.
5. Verify the following information and click **OK**:
   - Virtual Adapter
   - Backing Device
   - Remote Partition
   - Remote Adapter
   - Remote Backing Device

   **Note:** The virtual SCSI adapter fields may be blank if the mobile partition is powered off or if the physical disk has not been linked to the Virtual I/O Server's virtual SCSI adapter.
   If the information is incorrect, return to "Preparing the storage configuration for Partition Mobility" on page 234 and complete the task associated with the incorrect information.

## Preparing the network configuration for Partition Mobility

There are several tasks that you must complete to ensure your network configuration meets the minimal configuration for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

To prepare your network configuration for partition mobility, complete the following tasks.

*Table 65. Planning tasks for the network*

| Network planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Create a Shared Ethernet Adapter on the source and destination Virtual I/O Server logical partition using the HMC. See Creating a Shared Ethernet Adapter. | X | X |

*Table 65. Planning tasks for the network (continued)*

| Network planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 2. Configure virtual Ethernet adapters on the source and destination Virtual I/O Server logical partitions. | X | X |
| 3. Ensure that the mobile partition has a virtual Ethernet adapter. | X | |
| 4. Activate the mobile partition to establish communication between the virtual Ethernet adapter and Virtual I/O Server virtual Ethernet adapter. See Activating a partition profile for instructions. | X | |
| 5. Verify that the operating system of the mobile partition recognizes the new Ethernet adapter. | X | |
| 6. If any physical network adapters had been defined in the mobile partition, dynamically move the physical I/O using the HMC. See Managing physical I/O devices and slots dynamically using the HMC. | X | |
| 7. Set up the LAN so that the mobile partition can continue to communicate with other necessary clients and servers after the migration is completed. | X | X |

**Creating a Shared Ethernet Adapter using HMC version 7:**

You can create a Shared Ethernet Adapter on the Virtual I/O Server so that client logical partitions can access the external network without needing to own a physical Ethernet adapter.

If you plan to use a Shared Ethernet Adapter with a Host Ethernet Adapter (or Integrated Virtual Ethernet), ensure that the Logical Host Ethernet Adapter (LHEA) on the Virtual I/O Server is set to promiscuous mode. For instructions, see “Setting the LHEA to promiscuous mode” on page 85.

To create a Shared Ethernet Adapter on the Virtual I/O Server using the Hardware Management Console (HMC), version 7 or later, complete the following steps:

1. In the navigation area, expand **Systems Management** → **Servers** and select the server on which the Virtual I/O Server logical partition is located.
2. In the contents are, select the Virtual I/O Server logical partition.
3. Click **Tasks** and select **Configuration** → **Manage Profiles**. The Managed Profiles page is displayed.
4. Select the profile in which you want to create the Shared Ethernet Adapter and click **Actions** → **Edit**. The Logical Partition Profile Properties page is displayed.
5. Click the **Virtual Adapters** tab.
6. Click **Actions** → **Create** → **Ethernet adapter**.
7. Select **IEEE 802.1Q-compatible adapter**.
8. If you are using multiple VLANs, add any additional VLAN IDs for the client logical partitions that must communicate with the external network using this virtual adapter.
9. Select **Access external network** to use this adapter as a gateway between VLANs and an external network. This Ethernet adapter is configured as part of the Shared Ethernet Adapter.

10. If you are not using Shared Ethernet Adapter failover, you can use the default trunk priority. If you are using Shared Ethernet Adapter failover, then set the trunk priority for the primary share Ethernet adapter to a lower number than that of the backup Shared Ethernet Adapter.
11. When you are finished, click **OK**.
12. Assign or create one of the following real adapters:
    - Assign a physical Ethernet adapter to the Virtual I/O Server.
    - If you plan to aggregate more than one physical Ethernet adapter into a Link Aggregation or EtherChannel device, then assign multiple physical Ethernet adapters to the Virtual I/O Server.
    - If you plan to use the Shared Ethernet Adapter with a Host Ethernet Adapter, then create an LHEA for the Virtual I/O Server logical partition.
13. Click **OK** to exit the Logical Partition Profile Properties page.
14. Click **Close** to exit the Managed Profiles page.
15. Repeat this procedure for additional Shared Ethernet Adapters that you require.

When you are finished, configure the Shared Ethernet Adapter using the Virtual I/O Server command-line interface. For instructions, see “Configuring a Shared Ethernet Adapter” on page 86.

**Configuring a Shared Ethernet Adapter:**

Find instructions for configuring Shared Ethernet Adapters.

Before you can configure a Shared Ethernet Adapter, you must first create the adapter using the Hardware Management Console (HMC). For instructions, see one of the following tasks:
- “Creating a Shared Ethernet Adapter using HMC version 7” on page 84
- “Creating a Shared Ethernet Adapter using HMC version 6” on page 85

To configure a Shared Ethernet Adapter using the Virtual I/O Server, complete the following steps:
1. Verify that the virtual Ethernet trunk adapter is available by running the following command:
   ```
   lsdev -virtual
   ```
2. Identify the appropriate physical Ethernet adapter that will be used to create the Shared Ethernet Adapter by running the following command:
   ```
   lsdev -type adapter
   ```

   **Notes:**
   - Ensure that TCP/IP is not configured on the interface for the physical Ethernet adapter. If TCP/IP is configured, the mkvdev command in the next step fails.
   - You can also use a Link Aggregation, or EtherChannel, device as the Shared Ethernet Adapter.
   - If you plan to use the Host Ethernet Adapter or Integrated Virtual Ethernet with the Shared Ethernet Adapter, ensure that you use the Logical Host Ethernet Adapter to create the Shared Ethernet Adapter.
3. Configure the Shared Ethernet Adapter by running the following command:
   ```
   mkvdev -sea target_device -vadapter virtual_ethernet_adapters \
   -default DefaultVirtualEthernetAdapter -defaultid SEADefaultPVID
   ```

   Where:

   *target_device*

   The physical adapter being used as part of the Shared Ethernet Adapter device.

   *virtual_ethernet_adapters*

   The virtual Ethernet adapter or adapters that will use the Shared Ethernet Adapter.

*DefaultVirtualEthernetAdapter*
The default virtual Ethernet adapter used to handle untagged packets. If you have only one virtual Ethernet adapter for this logical partition, use it as the default.

*SEADefaultPVID*
The PVID associated with your default virtual Ethernet adapter.

For example, to create Shared Ethernet Adapter `ent3` with `ent0` as the physical Ethernet adapter (or Link Aggregation) and `ent2` as the only virtual Ethernet adapter (defined with a PVID of 1), type the following command:

```
mkvdev -sea ent0 -vadapter ent2 -default ent2 -defaultid 1
```

4. Verify that the Shared Ethernet Adapter was created by running the following command:

   ```
   lsdev -virtual
   ```

5. Do you plan to access the Virtual I/O Server from the network with the physical device used to create the Shared Ethernet Adapter?
   - Yes: Go to step 6 on page 87.
   - No: You are finished with this procedure and do not need to complete the remaining steps.
6. Do you plan to define IP addresses on any VLANs other than the VLAN specified by the PVID of the Shared Ethernet Adapter?
   - Yes: Go to step 7 on page 87 to create VLAN pseudo-devices.
   - No: Go to step 8 on page 87 to configure a TCP/IP connection.
7. To configure VLAN pseudo-devices, complete the following steps:
   a. Create a VLAN pseudo-device on the Shared Ethernet Adapter by running the following command:

      ```
      mkvdev -vlan TargetAdapter -tagid TagID
      ```

      Where:
      - *TargetAdapter* is the Shared Ethernet Adapter.
      - *TagID* is the VLAN ID that you defined when creating the virtual Ethernet adapter associated with the Shared Ethernet Adapter.

      For example, to create a VLAN pseudo-device using the Shared Ethernet Adapter `ent3` that you just created with a VLAN ID of 1, type the following command:

      ```
      mkvdev -vlan ent3 -tagid 1
      ```

   b. Verify that the VLAN pseudo-device was created by running the following command:

      ```
      lsdev -virtual
      ```

   c. Repeat this step for any additional VLAN pseudo-devices that you need.
8. Run the following command to configure the first TCP/IP connection. The first connection must be on the same VLAN and logical subnet as the default gateway.

   ```
   mktcpip -hostname Hostname -inetaddr Address -interface Interface -netmask \
   SubnetMask -gateway Gateway -nsrvaddr NameServerAddress -nsrvdomain Domain
   ```

   Where:
   - *Hostname* is the host name of the Virtual I/O Server
   - *Address* is the IP address you want to use for the TCP/IP connection
   - *Interface* is the interface associated with either the Shared Ethernet Adapter device or a VLAN pseudo-device. For example, if the Shared Ethernet Adapter device is ent3, the associated interface is en3.
   - *Subnetmask* is the subnet mask address for your subnet.
   - *Gateway* is the gateway address for your subnet.
   - *NameServerAddress* is the address of your domain name server.

- *Domain* is the name of your domain.

  If you do not have additional VLANs, then you are finished with this procedure and do not need to complete the remaining step.

9. Run the following command to configure additional TCP/IP connections:

   ```
   chdev -dev interface -perm -attr netaddr=IPaddress -attr netmask=netmask
   -attr state=up
   ```

   When using this command, enter the interface (en*X*) associated with either the Shared Ethernet Adapter device or VLAN pseudo-device.

The Shared Ethernet Adapter is now configured. After you configure the TCP/IP connections for the virtual adapters on the client logical partitions using the client logical partitions' operating systems, those logical partitions can communicate with the external network.

**Related concepts**

"Shared Ethernet Adapter failover" on page 58
Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server logical partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption.

## Validating the Partition Mobility environment using the HMC

This procedure provides instructions that explain how to validate the Partition Mobility environment using the HMC.

You must be a super administrator to perform this task.

To verify the Partition Mobility environment using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management**.
2. Select **Servers**.
3. In the navigation area, select the source server.
4. Select the mobile partition and expand **Operations > Mobility > Validate**. The Partition Migration Validation window opens.
5. In the validation panel, select the destination server and select **Validate**. The panel will be populated with suggested validation data.
6. Review the available virtual SCSI and virtual LAN settings on the destination system.
7. You may change the default settings of the mover service partition This is only valid for active Partition Mobility.
8. Click **Validate** to confirm that the changed settings are still acceptable for Partition Mobility.

## Migrating a logical partition using the Hardware Management Console

There are several tasks that you must complete to migrate a logical partition. Use this information to understand what you need to do to ensure that your migration is successful.

Prerequisites

To migrate a logical partition, complete the following tasks.

*Table 66. Prerequisite tasks for migrating a logical partition*

| Partition Mobility prerequisite tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. The PowerVM Enterprise Edition (or Advanced POWER Virtualization Enterprise Edition) hardware feature must be purchased and activated to use Partition Mobility. For more information about PowerVM Enterprise Edition, see PowerVM Editions. For instructions about activating the PowerVM Enterprise Edition hardware feature, see Entering the activation code for PowerVM Editions using the HMC version 7. | X | X |
| 2. Verify that you have completed all of the required planning tasks for Partition Mobility. See "Preparing for an HMC migration" on page 223 for instructions. | X | X |
| 3. Verify that the source and destination servers are in the Operating state. For instructions to power on the server, see Powering on a managed system. | X | X |
| 4. Verify that the mobile partition is powered off.<br>For instructions to power off the mobile partition, see the following information:<br>• AIX: Using the Hardware Management Console to shut down AIX logical partitions in the *Logical Partitioning Guide*[1].<br>• Linux: Using the Hardware Management Console to shut down Linux logical partitions in the *Logical Partitioning Guide*[1].<br>**Note:** If you want to actively move the logical partition, and the logical partition is in a crashed or failed state, then you need to return the logical partition to an operating state. See the following information for troubleshooting tips:<br>• AIX: Troubleshooting AIX logical partitions in the *Logical Partitioning Guide*[1].<br>• Linux: Troubleshooting Linux logical partitions in the *Logical Partitioning Guide*[1]. | | X |
| 5. Verify that the mobile partition is in the Operating state. For instructions to activate a logical partition, see Activating a partition profile. | X | |
| 6. Verify that the source and destination Virtual I/O Server logical partitions are active. For instructions to activate a logical partition, see Activating a partition profile. | X | X |
| 7. Verify that all tape and CD jobs are completed or stopped. | X | |

*Table 66. Prerequisite tasks for migrating a logical partition (continued)*

| **Partition Mobility prerequisite tasks** | **Active mobility task** | **Inactive mobility task** |
|---|---|---|
| 8. Run the migration verification tool on the HMC to verify that the servers, Virtual I/O Servers, mobile partition, storage, and network are ready for Partition Mobility. See "Validating the Partition Mobility environment using the HMC" on page 249 for instructions. | X | X |
| **Note:**<br>1. To view the PDF file of the *Logical Partitioning Guide* (SA76-0098), approximately 3 MB in size, see sa76-0098.pdf . | | |

To migrate a logical partition using the HMC, complete the following tasks:

1. In the navigation area, open **Systems Management**.
2. Select **Servers**.
3. In the contents area, open the source server.
4. Select the mobile partition and select **Operations > Mobility > Migrate**.
5. Complete the wizard.

Postrequisites

You must complete several postrequisite tasks after you migrate your logical partition. Use this information to understand what you need to do to ensure that your migration is successful.

*Table 67. Postrequisite tasks for migrating a logical partition*

| **Partition Mobility postrequisite tasks** | **Active mobility task** | **Inactive mobility task** |
|---|---|---|
| 1. Activate the mobile partition on the destination server. See Activating a partition profile for instructions. | | X |
| 2. (Optional) Add dedicated I/O adapters to the mobile partition on the destination server. See "Adding dedicated I/O to the mobile partition" on page 254 for instructions. | X | X |
| 3. If any virtual terminal connections were lost during the migration, re-establish the connections on the destination server. | X | X |
| 4. (Optional) Assign the mobile partition to a logical partition workload group. See "Adding the mobile partition to a partition workload group" on page 255 for instructions. | X | X |
| 5. If mobility-unaware applications were terminated on the mobile partition prior to its movement, then restart those applications on the destination. | X | |

## Powering on a managed system

Understand how to power on a managed system using the HMC.

You can use the HMC to power on a managed system and to monitor the power-on state.

To power on a managed system, you must be a member of one of the following roles:

- super administrator
- service representative
- operator
- product engineer

To power on a managed system, complete the following steps:

1. In the Navigation area, expand the **Systems Management** folder.
2. Click the **Servers** icon.
3. In the Contents area, select the managed system.
4. Select **Tasks**, then **Operations**, and then **Power On**
5. Select the desired power-on mode and click **OK**.

## Shutting down AIX logical partitions using the HMC

You can shut down AIX logical partitions using the Hardware Management Console (HMC).

**Delayed shutdown of the operating system:**

When you use the delayed shutdown option, the HMC issues the AIX **shutdown** command to shut down the logical partition normally. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.

To perform a delayed shutdown of the operating system using the HMC, complete the following:

1. In the navigation area of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition is located.
2. In the contents area, select the logical partition, click the **Tasks** button, and choose **Operations** → **Shut Down**.
3. Select **Operating System** and click **OK**.

**Immediate shutdown of the operating system:**

When you use the immediate shutdown option, the HMC issues the AIX **shutdown -F** command to shut down the logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.

To perform an immediate shutdown of the operating system using the HMC, complete the following:

1. In the navigation area of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition is located.
2. In the contents area, select the logical partition, click the **Tasks** button, and choose **Operations** → **Shut Down**.
3. Select **Operating System Immediate** and click **OK**.

**Delayed shutdown of a logical partition:**

When you use the delayed shutdown option, the logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks.

To perform a delayed shutdown of an AIX logical partition using the HMC, complete the following:

1. In the navigation area of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition is located.
2. In the contents area, select the logical partition, click the **Tasks** button, and choose **Operations** → **Shut Down**.
3. Select **Delayed** and click **OK**.

**Immediate shutdown of a logical partition:**

When you use the immediate shutdown option, the system shuts down without any preset delay.

To perform an immediate shutdown of an AIX logical partition using the HMC, complete the following:

1. In the navigation area of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition is located.
2. In the contents area, select the logical partition, click the **Tasks** button, and choose **Operations** → **Shut Down**.
3. Select **Immediate** and click **OK**.

## Shutting down Linux logical partitions using the HMC

You can shut down Linux logical partitions and the Linux operating system using the Hardware Management Console (HMC).

**Delayed shutdown of the operating system:**

When you use the delayed shutdown option, the Hardware Management Console (HMC) issues the Linux **shutdown -h +1** command to shut down the logical partition normally. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.

To perform a delayed shutdown of the operating system using the HMC, complete the following:

1. In the navigation area of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition resides.
2. In the contents area, select the logical partition, click the **Tasks** button, and choose **Operations** → **Shut Down**.
3. Select **Operating System** and click **OK**.

**Immediate shutdown of the operating system:**

When you use the immediate shutdown option, the Hardware Management Console (HMC) issues the Linux **shutdown -h now** command to shut down the logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.

To perform an immediate shutdown of the operating system using the HMC, complete the following:

1. In the navigation area of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition resides.
2. In the contents area, select the logical partition, click the **Tasks** button, and choose **Operations** → **Shut Down**.
3. Select **Operating System Immediate** and click **OK**.

**Delayed shutdown of a logical partition:**

When you use the delayed shutdown option, the logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks.

To perform a delayed shutdown of a Linux logical partition using the HMC, complete the following:

1. In the navigation area of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition resides.
2. In the contents area, select the logical partition, click the **Tasks** button, and choose **Operations** → **Shut Down**.
3. Select **Delayed** and click **OK**.

**Immediate shutdown of a logical partition:**

When you use the immediate shutdown option, the system shuts down without any preset delay.

To perform an immediate shutdown of a Linux logical partition using the HMC, complete the following:

1. In the navigation area of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition resides.
2. In the contents area, select the logical partition, click the **Tasks** button, and choose **Operations** → **Shut Down**.
3. Select **Immediate** and click **OK**.

## Adding dedicated I/O to the mobile partition

This procedure provides post-migration instructions that explain how to add dedicated I/O to the mobile partition.

You must be a super administrator to complete this task.

Dedicated I/O allow your managed system to gather, store, and transmit data. Dedicated I/O devices are physical devices found in the server itself and in expansion units and towers that are attached to the server.

Prior to moving the mobile partition from the source environment to the destination environment, you might have moved dedicated I/O from the mobile partition to another logical partition. Now that you have successfully moved the mobile partition to the destination environment, you can add dedicated I/O to the mobile partition.

To add dedicated I/O to the mobile partition using the HMC, complete the following steps:

1. Dynamically add dedicated I/O adapters to the mobile partition, or dynamically move dedicated I/O adapters from a logical partition to the mobile partition. See the following instructions:
   - AIX: Dynamically managing physical I/O devices and slots
   - Linux: Adding physical I/O devices and slots dynamically using version 7 or later of the HMC
2. Update the partition profiles with the new I/O assignments: You must update the partition profiles so that the next time they are activated, the new I/O assignments are not lost.
   a. In the navigation area, open **Systems Management** and select **Servers**.
   b. Select the managed server of your choice in the navigation area.
   c. Select the logical partition of your choice in the contents area.
   d. Select **Configuration > Manage Profiles.**
   e. Select the profile of your choice and select **Actions > Edit**.
   f. Click the **I/O** tab.
   g. For each resource that you want to be dedicated to the mobile partition, select **Add as required** or **Add as desired**.

   h. Click **OK**. In order for this change to take affect, you will need to activate this logical partition with this profile.

   i. Repeat step 2f through 2h for each partition profile associated with the mobile partition.

**Managing physical I/O devices and slots dynamically using the HMC:**

You can add, remove, and move physical I/O devices and slots dynamically to and from running logical partitions using the Hardware Management Console (HMC). This allows logical partitions to share infrequently used I/O devices (such as optical disk drives).

Logical partitions can have desired or required I/O devices or slots. When you specify that an I/O device or slot is desired (or shared), this means either that the I/O device or slot is meant to be shared with other logical partitions, or that the I/O device or slot is optional. When you specify that an I/O device or slot is required (or dedicated), then you cannot activate the logical partition if the I/O device or slot is unavailable or in use by another logical partition.

**Note:** If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new partition configuration, either change the partition profile or save the partition configuration to a new partition profile.

**Adding physical I/O devices and slots dynamically using the HMC:**

You can add a physical I/O slot (and the adapter and devices that are connected to that slot) to a running logical partition using the Hardware Management Console (HMC). This allows you to add I/O capabilities to a running logical partition without having to shut down the logical partition.

A Linux logical partition supports the dynamic addition of physical I/O slots only if the following conditions are met:

- A Linux distribution that supports dynamic logical partitioning is installed on the Linux logical partition. Distributions that support dynamic logical partitioning include SUSE Linux Enterprise Server 9 and later versions.
- The DynamicRM tool package is installed on the Linux logical partition. For more information on the DynamicRM tool package, see the Service and productivity tools Web site.

To add a physical I/O slot dynamically to a running logical partition using the HMC, you must be a super administrator, service representative, product engineer, or operator. For more information about user roles, refer to Tasks and roles in the *Operations Guide for the Hardware Management Console and Managed Systems*. To view the abstract of *Operations Guide for the Hardware Management Console and Managed Systems* (SA76-0085), see sa76-0085.pdf .

To add a physical I/O slot dynamically to a running logical partition using the HMC, follow these steps:

1. In the navigation pane of your HMC, open **Systems Management**, open **Servers**, and click the managed system on which the logical partition is located.
2. In the contents pane, select the logical partition, click the **Tasks** button, and choose **Dynamic Logical Partitioning** → **Physical Adapters** → **Add**.
3. Select the physical I/O slot that you want to add to the logical partition.
4. If you want to assign the physical I/O slot to an I/O pool, select the I/O pool for the physical I/O slot in **I/O Pool ID**.
5. Click **OK**.

## Adding the mobile partition to a partition workload group

This procedure provides post-migration instructions that explain how to add the mobile partition to a partition workload group.

You must be a super administrator to complete this task.

A partition workload group identifies a set of logical partitions that are located on the same physical system. Workload management tools use partition workload groups to identify which logical partitions they can manage.

Prior to moving the mobile partition from the source environment to the destination environment, you might have removed the mobile partition from a partition workload group. Now that you have successfully moved the mobile partition to the destination environment, you can add it to a partition workload group.

To add the mobile partition to a partition workload group using the HMC, complete the following steps:

1. In the navigation area, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation area.
3. Select the logical partition of your choice in the contents area.
4. Select **Configuration > Manage Profiles**.
5. Select the profile of your choice and select **Actions > Edit.**.
6. Click the **Settings** tab.
7. In the Workload Management area, select (None) and click **OK**.
8. Repeat steps 1 through 7 for all partition profiles associated with the mobile partition. In order for this change to take affect, you will need to activate this logical partition with this profile.

This can also be changed using DLPAR by selecting the logical partition > **Properties > Other** tab.

# Troubleshooting problems with Partition Mobility

Learn how to understand, isolate, and resolve problems related to active and inactive Partition Mobility.

Sometimes you will be able to resolve a problem on your own, while at other times you will need to gather information to help the service technicians resolve your problem in a timely manner.

## Troubleshooting active Partition Mobility problems

Learn how to troubleshoot problems that might occur with active Partition Mobility.

The following is a list of possible errors and ways to recover.

*Table 68. Known problems and solutions for active Partition Mobility*

| Problem | Solution |
|---|---|
| If the operating system running in the mobile partition does not explicitly support the processor version register of the destination server, and the processor determines that explicit support is required, then the processor will not allow the migration to proceed. | Either:<br>• Move the logical partition to another system.<br>• Update the operating system to a level that supports the target system processor version registers. |
| You receive an error concerning the operating system while attempting to migrate a logical partition. | 1. Examine the operating system error logs for operating system-related failures.<br>2. Examine the HMC log for application-related failures. |
| You receive an HMC error concerning insufficient memory on the destination server.<br>**Note:** Sufficient memory includes the amount of available memory on the server and the amount of available contiguous memory on the server. If the mobile partition requires more contiguous memory, making more memory available will not solve the problem. | Either:<br>• Move the logical partition to a different server.<br>• Make more memory available on the destination server. See “Determining available memory on the destination server” on page 224 for instructions. |

*Table 68. Known problems and solutions for active Partition Mobility (continued)*

| Problem | Solution |
|---|---|
| The HMC and managed systems lost their connection while the migration was in progress. | 1. In the HMC navigation area, select **Partition Migration**.<br>2. In the contents area, select either the source or destination system and select **Recovery**.<br>3. From the Migration Recovery panel, click **Recover**. |
| While attempting to change resources dynamically, you receive an error that the Resource Monitoring and Control (RMC) daemon is not connected. | This error typically occurs when there is a network connection problem between the logical partitions and the HMC. To resolve this error, check your system network setup. |

### Troubleshooting inactive Partition Mobility problems

Learn how to troubleshoot problems with inactive Partition Mobility.

The following is a list of possible errors and ways to recover.

*Table 69. Known problems and solutions for inactive Partition Mobility*

| Problem | Solution |
|---|---|
| If the mobile partition is moved to a server that the operating system does not support (and explicit support is required), then the boot of the logical partition on the destination server will fail. | Move the logical partition to a different system. |
| You receive an error concerning insufficient memory on the destination server.<br>**Note:** Sufficient memory is the amount of available memory the available contiguous memory on the destination server. Thus, if the mobile partition requires more contiguous memory, making more memory available will not solve the problem. | Either:<br>• Move the logical partition to a different server.<br>• Make more memory available on the destination server. See “Determining available memory on the destination server” on page 224 for instructions. |

### Reference codes for Partition Mobility

Partition Mobility reference codes are diagnostic aids that help you determine the source of a hardware or operating system problem. Using reference codes, you can find the correct solution to fix the problem. To use reference codes effectively, you must use them in conjunction with other service and support procedures.

To help you better understand your problem and how you can fix it, refer to the System Reference Codes PDF for your server model.

# Using the Integrated Virtualization Manager for Live Partition Mobility

Learn more about using the Integrated Virtualization Manager to migrate an active or inactive logical partition.

## Integrated Virtualization Manager environment

Use this information to help gain an understanding of an active or inactive logical partition environment using the Integrated Virtualization Manager.

### Source and destination servers

This information describes how to set up the environment for the source and destination servers before performing a migration.

Two servers are involved in Partition Mobility. The *source server* is the server from which you want to move the logical partition, and the *destination server* is the server to which you want to move the logical partition. The source and destination servers must be POWER6 processor-based servers to participate in Partition Mobility. The destination server must have enough available processor and memory resources to allow the mobile partition to run on its server.

### Integrated Virtualization Manager

When you install the Virtual I/O Server on a system that is not managed by an HMC or an BladeCenter blade server, the Virtual I/O Server becomes the management partition and provides the Integrated Virtualization Manager for systems management. The Integrated Virtualization Manager provides a Web-based and command-line interface that you can use to migrate a logical partition from one POWER6 processor-based system to another.

The Migration task on the Integrated Virtualization Manager helps you validate and complete a partition migration. The Integrated Virtualization Manager determines the appropriate type of migration to use based on the state of the logical partition. If the logical partition is in the *Running* state, then the migration is active. If the logical partition is in the *Not Activated* state, then the migration is inactive. Before migrating your logical partition, conduct a validation check to ensure your migration will complete successfully.

### Networking

During active Partition Mobility, it is important that the two management partitions be able to communicate with each other. The network is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access. The virtual LAN must be bridged to a physical network using a virtual Ethernet bridge in the management partition. The LAN must be configured so that the mobile partition can continue to communicate with other necessary clients and servers after a migration is completed.

Active Partition Mobility has no specific requirements on the mobile partition's memory size. The memory transfer is a procedure that does not interrupt a mobile partition's activity and may take time when a large memory configuration is involved on a slow network. Because of this, use a high-bandwidth connection, such as Gigabit Ethernet.

### Storage configuration for partition migration

This topic summarizes the storage configuration required for partition mobility.

The mobile partition moves from one server to another by the source server sending the logical partition state information to the destination server over a local area network. However, partition disk data cannot pass from one system to another system over a network. Thus, for partition mobility to succeed, the mobile partition must use storage resources virtualized by a storage area network so that it can access the same storage from both the source and destination servers.

## Requirements for Partition Mobility using the Integrated Virtualization Manager

Learn more about the general software and hardware requirements for Partition Mobility using the Integrated Virtualization Manager.

The hardware and software required to use Partition Mobility varies depending on whether you are migrating an active or inactive AIX or Linux logical partition. Make sure that your Partition Mobility environment meets minimum requirements before you migrate your logical partition.

## Source and destination server requirements

The following table shows the hardware requirements for the Integrated Virtualization Manager.

*Table 70. Source and destination server requirements*

| Server requirements | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The source and destination server must be one of the following POWER6 models:<br>• 03E/4A<br>• 04E/8A<br>• JS/12 Express<br>• JS/22 Express | X | X |

## Integrated Virtualization Manager requirements

The following table shows the software requirements for the Integrated Virtualization Manager.

*Table 71. Integrated Virtualization Manager software requirements*

| Integrated Virtualization Manager requirement | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The source and destination servers must be using the Integrated Virtualization Manager at version 1.5 or later. For instructions about how to install the Integrated Virtualization Manager, see Installing the Integrated Virtualization Manager.<br>To determine the current Integrated Virtualization Manager version and update it if necessary, see Viewing and updating the code level of the management partition. | X | X |
| The PowerVM Enterprise Edition (or Advanced POWER Virtualization Enterprise Edition) hardware feature must be purchased and activated to use Partition Mobility. For more information about PowerVM Enterprise Edition, see PowerVM Editions. For instructions about activating the PowerVM Enterprise Edition hardware feature, see Entering the activation code for PowerVM Editions with the Integrated Virtualization Manager. | X | X |

## Operating system requirements

The following table shows the supported software requirements for Partition Mobility.

*Table 72. Operating system requirements*

| Operating system requirement | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| The operating system running in the mobile partition must be AIX or Linux. | X | X |
| The operating system must be at one of the following levels:<br>• AIX 5L Version 5.3 with the 5300-07 Technology Level or later<br>• SUSE Linux Enterprise Server 10 (SLES 10) Service Pack 1 or later<br>Earlier versions of AIX and Linux can participate in inactive Partition Mobility if the operating systems support virtual devices and POWER6 models. | X | X |

## Storage requirements

The following table shows the storage requirements for the source and destination server.

*Table 73. Source and destination server storage requirements*

| Storage requirements | Active mobility requirement | Inactive mobility requirement |
|---|---|---|
| Storage Area Network (SAN) is the only supported storage for Partition Mobility. | X | X |
| The mobile partition must be using storage that is visible to the Virtual I/O Servers on both the source and destination systems. | X | X |

# Preparing for an Integrated Virtualization Manager migration

Use this information to help gain an understanding of what to consider when planning to migrate an active or inactive logical partition using the Integrated Virtualization Manager.

## Preparing the source and destination servers for an Integrated Virtualization Manager partition migration

Complete the tasks to prepare the source and destination servers for an Integrated Virtualization Manager partition migration.

There are several tasks that you must complete to prepare the source and destination server an for an Integrated Virtualization Manager partition migration. Use this information to understand what you need to do to ensure that your migration is successful.

To prepare the source and destination server for an Integrated Virtualization Manager partition migration, complete the following tasks.

*Table 74. Planning tasks for the source and destination servers*

| Server planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Ensure that the source and destination servers meet the requirements for Partition Mobility using theIntegrated Virtualization Manager . See "Requirements for Partition Mobility using the Integrated Virtualization Manager" on page 258 for information. | X | X |
| 2. Ensure that the logical memory block size is the same on the source and destination server. Determine the logical memory block size of each server, and update the sizes if necessary. See Viewing and modifying system properties for instructions.<br>**Note:** The Validate function checks this for you. | X | X |
| 3. Ensure that the destination server has enough available memory to support the mobile partition. See "Determining available memory on the destination server using the Integrated Virtualization Manager" for instructions.<br>**Note:** The Validate function checks this for you. | X | X |
| 4. Ensure that the destination server has enough available processors to support the mobile partition. See Determining available processors on the destination server using the Integrated Virtualization Manager for instructions.<br>**Note:** The Validate function checks this for you. | X | X |
| 5 Verify that the source and destination Virtual I/O Server can communicate with each other.<br>**Note:** The Validate function checks this for you. | X | X |

### Determining available memory on the destination server using the Integrated Virtualization Manager:

This procedure provides instructions that explain how to determine the available memory on the destination server and allocate more memory if necessary using the Integrated Virtualization Manager.

Use any role other than View Only to perform this task. Users with the Service Representative (SR) user role cannot view or modify storage values.

To determine the available memory on the destination server using the Integrated Virtualization Manager, complete the following steps:

1. Determine how much memory the mobile partition requires:

a. From the Partition Management menu, click **View/Modify Partition**. The View/Modify Partition panel is displayed.
b. Select the logical partition for which you want to view the properties.
c. From the Tasks menu, click **Properties**. The Partition Properties panel is displayed.
d. Click the **Memory** tab and record the minimum, maximum, and available memory settings.
e. Click **OK**

2. Determine the memory available on the destination server:
   a. From the Partition Management menu, click **View/Modify System Properties**. The View/Modify System Properties panel is displayed.
   b. Select the**Memory** tab.
   c. Record the **Current memory available** .
   d. Click **Apply**.
3. Compare the values from steps 1 and 2.
   - If the destination server has enough available memory to support the mobile partition, continue with "Preparing the source and destination servers for an Integrated Virtualization Manager partition migration" on page 260.
   - If the destination server does not have enough available memory to support the mobile partition, use the Integrated Virtualization Manager to dynamically remove memory from the logical partition or you can remove memory from logical partitions on the destination server.

**Note:** The POWER6 hypervisor uses some memory on the destination server after the migration. Ensure you have enough memory on the destination server to support the logical partition and the hypervisor requirements.

**Determining available processors on the destination server using the Integrated Virtualization Manager:**

This procedure provides instructions that explain how to determine the available processors on the destination server and allocate more processors if necessary.

You must be a super administrator to perform this task.

To determine the available processors on the destination server using the using the Integrated Virtualization Manager, complete the following steps:

1. Determine how many processors the mobile partition requires:
   a. From the Partition Management menu, click **View/Modify Partition**. The View/Modify Partition panel is displayed.
   b. Select the logical partition for which you want to view the properties.
   c. From the Tasks menu, click **Properties**. The Partition Properties panel is displayed.
   d. Click the **Processing** tab and record the minimum, maximum, and available processing units settings.
   e. Click **OK**
2. Determine the processors available on the destination server:
   a. From the **Partition Management** menu, click **View/Modify System Properties**. The View/Modify System Properties panel is displayed.
   b. Select the**Processing** tab.
   c. Record the **Current processing units available**.
   d. Click **Apply**.
3. Compare the values from steps 1 and 2.

- If the destination server has enough available processors to support the mobile partition, then continue with "Preparing the source and destination servers for an Integrated Virtualization Manager partition migration" on page 260.
- If the destination server does not have enough available processors to support the mobile partition, use the Integrated Virtualization Manager to dynamically remove the processors from the logical partition or you can remove processors from logical partitions on the destination server.

## Preparing the Integrated Virtualization Manager for Partition Mobility

To prepare the Integrated Virtualization Manager for Partition Mobility, complete the following task.

*Table 75. Planning tasks for the Integrated Virtualization Manager*

| Integrated Virtualization Manager planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| Ensure that the Integrated Virtualization Manager meets the requirements for Partition Mobility. See "Requirements for Partition Mobility using the Integrated Virtualization Manager" on page 258 for information. | X | X |

## Preparing the mobile partition for Partition Mobility using the Integrated Virtualization Manager

There are several tasks that you must complete to prepare the mobile partition for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

To prepare the mobile partition for Partition Mobility using the Integrated Virtualization Manager, complete the following tasks.

*Table 76. Planning tasks for the mobile partition*

| Mobile partition planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Ensure that the operating system meets the requirements for Partition Mobility. See "Requirements for Partition Mobility using the Integrated Virtualization Manager" on page 258 for information. | X | |
| 2. Ensure that the source and destination management partitions can communicate to each other. | X | X |
| 3. Ensure that the mobile partition is not part of a partition workload group. See "Removing the mobile partition from a partition workload group using the Integrated Virtualization Manager" on page 264 for instructions. | X | X |

*Table 76. Planning tasks for the mobile partition (continued)*

| Mobile partition planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 4. Ensure that the mobile partition does not have physical adapters or a Host Ethernet Adapter (sometimes referred to as Integrated Virtual Ethernet). See Dynamically managing physical adapters and Assigning Host Ethernet Adapter port to a logical partition for instructions.<br>**Note:** The Integrated Virtualization Manager will remove any physical I/O assigned to the mobile partition during an inactive migration. | X | |
| 5. Ensure that the applications running in the mobile partition are mobility-safe or mobility-aware. See Software applications that recognize migrations for more information. | X | |

**Removing the mobile partition from a partition workload group using the Integrated Virtualization Manager:**

This procedure provides instructions that explain how to remove the mobile partition from a partition workload group.

A partition workload group identifies a set of logical partitions that are located on the same physical system. A partition workload group is defined when you use the Integrated Virtualization Manager to configure a logical partition. The partition workload group is intended for applications that manage software groups. For a logical partition to participate in Partition Mobility, it cannot be assigned to a partition workload group.

To remove the mobile partition from a partition workload group using the Integrated Virtualization Manager, complete the following steps:

1. From the Partition Management menu, click **View/Modify Partition**. The View/Modify Partition window is shown.
2. Select the logical partition that you want to remove from the partition workload group.
3. From the Tasks menu, and click **Properties**. The Partition Properties window is shown.
4. In the General tab, deselect **Partition workload group participant**.
5. Click **OK**.

**Removing physical adapters from the mobile partition:**

This procedure provides instructions that explain how to remove physical adapters from the mobile partition.

Physical adapters allow your managed system to gather, store, and transmit data. Physical adapters are physical devices found in the server unit itself and in expansion units and towers that are attached to the server.

For a logical partition to participate in active Partition Mobility, it cannot have physical adapters. All adapters must be virtual. You can use inactive Partition Mobility if a mobile partition has physical adapters. The adapters will automatically be removed from the logical partition before the migration occurs.

To remove physical adapters from the mobile partition using the Integrated Virtualization Manager, complete the following steps:

1. From the Partition Management menu, click **View/Modify Partition**. The View/Modify Partition panel is shown.
2. Select the logical partition that you want to change.
3. From the Tasks menu, and click **Properties**. The Partition Properties panel is shown.
4. Click the **Physical Adapters** tab
5. Deselect the physical adapters that you would like to remove from the logical partition.
6. Click **OK**.

## Preparing the storage configuration for Partition Mobility using the Integrated Virtualization Manager

There are several tasks that you must complete to ensure your storage configuration meets the minimal configuration for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

If your environment already meets the minimal configuration previously shown, then select verify Verifying the existing storage configuration.

If you need to set up this environment, then select Configuring the storage environment.

**Verifying the existing storage configuration for Partition Mobility using the Integrated Virtualization Manager:**

Complete the tasks to verify that your storage configuration meets the minimal configuration requirements for Partition Mobility.

Before you migrate your logical partition, you need to verify that your existing storage configuration meets the requirements needed to migrate your logical partition. Use the following tasks to understand how your storage configuration is set up.

Verify the following attributes, assignments, and connections.

*Table 77. Planning tasks for storage*

| Storage planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Verify that virtual devices can be migrated. See Identifying exportable disks | X | |
| 2. Verify that the mobile partition has access to the physical storage. See “Verifying that the mobile partition has access to its physical storage” on page 266 for instructions. | X | X |
| 3. Verify that the mobile partition does not have physical adapters or a Host Ethernet Adapter (or Integrated Virtual Ethernet). See “Removing physical adapters from the mobile partition” on page 264 and Assigning Host Ethernet Adapter port to a logical partition for instructions. | X | X |

**Configuring the storage environment for Partition Mobility:**

There are several tasks that you must complete to configure your storage environment for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

To set up your storage configuration, complete the following tasks.

*Table 78. Planning tasks for storage*

| Storage planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Set up and configure a storage area network. See Virtual I/O Server Supported Environment for a list of storage devices supported by the Virtual I/O Server. | X | X |
| 2. Connect the source and destination management partitions to the storage area network. | X | X |
| 3. Set the reserve_policy attributes on the mobile client physical volumes to no_reserve. | X | X |
| 4. Assign the same disks to the source and all possible destination management partitions. | X | X |
| 5. Verify that the mobile partition has access to the SAN device. | X | X |
| 6. Verify that the mobile partition does not have physical or dedicated I/O adapters and devices. See "Removing physical adapters from the mobile partition" on page 264 for instructions. | X | |

*Verifying that the mobile partition has access to its physical storage:*

This procedure provides instructions that explain how to verify that the mobile partition has access to the physical storage on the storage area network.

For Partition Mobility to be successful, the mobile partition must have access to the same physical storage from both the source and destination environments. In the destination environment, the SAN host-attached adapter on the destination management partition must be connected to the same storage area network as the source management partition and have access to the same mobile partition physical storage as the source management partition

To verify these connections using the Integrated Virtualization Manager, complete the following steps:

1. From the Virtual Storage Management menu, click **View/Modify Virtual Storage**.
2. On the Virtual Disk tab, verify that the logical partition does not own any virtual disk.
3. On the Physical Volumes tab, verify the physical volumes mapped to the mobile partition are exportable. See Identifying exportable disks for more information.

   If the information is incorrect, return to Preparing the storage configuration and complete the task associated with the incorrect information.

## Preparing the network configuration for Partition Mobility using the Integrated Virtualization Manager

You must complete several tasks to ensure your network configuration meets the minimal configuration for Partition Mobility. Use this information to understand what you need to do to ensure that your migration is successful.

To prepare your network configuration for Partition Mobility , complete the following tasks.

*Table 79. Planning tasks for the network*

| Network planning tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Configure a virtual Ethernet bridge on the source and destination management partition using the Integrated Virtualization Manager. See "Configuring virtual Ethernet bridges on the managed system using the Integrated Virtualization Manager" on page 179.<br>Next view the virtual Ethernet settings for the managed system and change the virtual Ethernet network configuration on the source and destination management partition. See "Configuring virtual Ethernet bridges on the managed system using the Integrated Virtualization Manager" on page 179 | X | X |
| 2. Ensure you connect the source and destination management partition and the shared Ethernet adapter to the network. | X | X |
| 3. Ensure that the mobile partition has a virtual Ethernet adapter. | X | |
| 4. Activate the mobile partition to establish communication between the virtual Ethernet and management partition virtual Ethernet adapter. See Activating logical partitions for instructions. | X | |
| 5. Verify that the operating system of the mobile partition recognizes the new Ethernet adapter. To configure and manage new Ethernet adapters, see Adapter management and configuration in the servers and AIX Information Center. | X | |
| 6. Dynamically remove the physical I/O using the Integrated Virtualization Manager. See Dynamically managing I/O . | X | |
| 7. Set up the LAN so that the migrating logical partition can continue to communicate with other necessary clients and servers after the migration is completed. | X | X |

## Validating the Partition Mobility environment using the Integrated Virtualization Manager

This procedure provides instructions that explain how to validate the Partition Mobility environment using the Integrated Virtualization Manager.

To validate the Partition Mobility environment using the Integrated Virtualization Manager, complete the following steps:

1. From the Partition Management menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition for which you want to migrate and from the Tasks menu, select **Migrate**.
3. Enter the **Remote IVM or HMC**, **Remote user ID**, and **Password** of the logical partition you plan to migrate.
4. Click **Validate** to confirm that the changed settings are acceptable for Partition Mobility.

## Migrating a logical partition using the Integrated Virtualization Manager

You must complete several tasks to migrate a logical partition. Use this information to understand what you need to do to ensure that your migration is successful.

Prerequisites

There are several prerequisite tasks that you must complete to migrate your logical partition. Use this information to understand what you need to do to ensure that your migration is successful.

*Table 80. Prerequisite tasks for migrating a logical partition*

| Partition Mobility prerequisite tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1.The PowerVM Enterprise Edition (or Advanced POWER Virtualization Enterprise Edition) hardware feature must be purchased and activated to use Partition Mobility. For more information about PowerVM Enterprise Edition, see PowerVM Editions. For instructions about activating the PowerVM Enterprise Edition hardware feature, see Entering the activation code for PowerVM Editions using the HMC version 7. | X | X |
| 2. Verify that you have completed all of the required planning tasks for Partition Mobility. See “Preparing for an Integrated Virtualization Manager migration” on page 260 for instructions. | X | X |
| 3. Verify that the memory and processor resources are synchronized after dynamically adding or removing resources. See “Dynamically managing memory” on page 185 and “Dynamically managing processing power” on page 186 for more information. | | |
| 4. Verify that the source and destination servers are in the Operating state. | X | X |
| 5. Verify that the mobile partition is powered off. | | X |
| 6. Verify that the mobile partition is in the Operating state. For instructions to activate a logical partition, see Activating a logical partition. | X | |

*Table 80. Prerequisite tasks for migrating a logical partition (continued)*

| Partition Mobility prerequisite tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 7. Verify that the source and destination Virtual I/O Servers are active. For instructions to activate a logical partition, see Activating a logical partition. | X | X |
| 8. Verify that all tape and CD jobs are completed or stopped. | X | |
| 9. Run the migration validation tool on the Integrated Virtualization Manager to verify that the servers, mobile partition, storage, and network are ready for Partition Mobility. See "Validating the Partition Mobility environment using the Integrated Virtualization Manager" on page 268 for instructions. | X | X |

## Migrating a logical partition using the Integrated Virtualization Manager

To migrate a logical partition using the Integrated Virtualization Manager, complete the following tasks:

1. From the Partition Management menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition for which you want to migrate and from the Tasks menu, select **Migrate**.
3. Enter the **Remote IVM**, **Remoter user ID**, and **Password** of the logical partition you plan to migrate.
4. Click **Migrate**.

## Postrequisites

There are several postrequisite tasks that you must complete to migrate your logical partition. Use this information to understand what you need to do to ensure that your migration is successful.

*Table 81. Postrequisite tasks for migrating a logical partition*

| Partition Mobility postrequisite tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 1. Activate the mobile partition on the destination server. See Activating a logical partition for instructions. | | X |
| 2. (Optional) Add physical adapters to the mobile partition on the destination server. See Dynamically managing physical adapters for instructions. | X | X |
| 3. If any virtual terminal connections were lost during the migration, re-establish the connections on the destination server. See Opening a virtual terminal session for instructions. | X | X |
| 4. (Optional) Assign the mobile partition to a logical partition group. See "Adding a client logical partition to the partition workload group" on page 183 for instructions. | X | X |

*Table 81. Postrequisite tasks for migrating a logical partition (continued)*

| Partition Mobility postrequisite tasks | Active mobility task | Inactive mobility task |
|---|---|---|
| 5. If mobility-unaware applications were terminated on the mobile partition prior to its movement, then restart those applications on the destination. | X | |

## Shutting down logical partitions

Use the Integrated Virtualization Manager to shut down the selected logical partitions or the entire managed system.

Use any role other than View Only to perform this task.

The Integrated Virtualization Manager provides the following types of shutdown options for logical partitions:

- Operating System (recommended)
- Delayed
- Immediate

The recommended shutdown method is to use the client operating systems shutdown command. Using the immediate shutdown method should be used as a last resort as this causes an abnormal shutdown which might result in data loss.

If you choose the Delayed shutdown method, then be aware of the following considerations:

- Shutting down the logical partitions is equivalent to pressing and holding the white control-panel power button on a server that is not partitioned.
- Use this procedure only if you cannot successfully shut down the logical partitions through operating system commands. When you use this procedure to shut down the selected logical partitions, the logical partitions wait a predetermined amount of time to shut down. This allows the logical partitions time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it ends abnormally, and the next restart might take a long time.

If you plan to shut down the entire managed system, shut down each client logical partition and then shut down the Virtual I/O Server management partition.

To shut down a logical partition, do the following:

1. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition that you want to shut down.
3. Click **Shutdown**. The Shutdown Partitions panel is displayed.
4. Select the shutdown type.
5. Optional: Select **Restart after shutdown completes** if you want the logical partition to start immediately after it shuts down.
6. Click **OK** to shut down the partition. The View/Modify Partitions panel is displayed, and the partition is shut down.

For more information about shutting down logical partitions, see the online help (?).

## Dynamically managing physical adapters

You can add and remove physical adapters to and from a running logical partition.

You can change the physical adapter settings for a logical partition at any time if the partition is capable of dynamic I/O adapter changes.

When making dynamic I/O adapter changes, keep the following items in mind:

- You might lose data if you remove a physical adapter from a running logical partition.
- You cannot assign a physical adapter to another partition if it is being used by the operating system of the partition to which it is currently assigned. If you attempt to reassign the adapter, an error message is displayed. You must unconfigure the device by using the tools of the appropriate operating system before you can change the adapter's partition assignment.

Before you start, ensure that the Integrated Virtualization Manager is at version 1.5 or later. To update the Integrated Virtualization Manager, see "Viewing and updating the code level of the Integrated Virtualization Manager management partition" on page 199.

To dynamically add or remove physical adapters to or from a running logical partition, follow these steps:

1. If no client logical partitions exist, go to step 4 on page 181.
2. Select the logical partition to which you want to assign a physical adapter and click **Properties**.
3. Verify that **Yes** is displayed for **I/O adapter DLPAR Capable**. You might need to click **Retrieve Capabilities** to verify this value. If **No** is displayed for **Processing DLPAR Capable**, then you cannot dynamically add or remove physical adapters to or from the logical partition.
4. From the **I/O Adapter Management** menu, click **View/Modify Physical Adapters**.
5. Select the adapter whose partition assignment you want to change and click **Modify Partition Assignment**.
6. Select the logical partition to which you want to assign the physical adapter and click **OK**. If you want to make this adapter available to any client logical partition, including those not yet created, select **None** as the **New partition**.

## Adding a client logical partition to the partition workload group

If you want to manage logical partition resources using a workload management tool, then you need to add the client logical partition to the partition workload group.

A *partition workload group* identifies a set of logical partitions that are located on the same physical system. Workload management tools use partition workload groups to identify which logical partitions they can manage. For example, Enterprise Workload Manager (EWLM) can dynamically and automatically redistribute processing capacity within a partition workload group to satisfy workload performance goals. EWLM adjusts processing capacity based on calculations that compare the actual performance of work processed by the partition workload group to the business goals defined for the work.

Workload management tools use dynamic logical partitioning (DLPAR) to make resource adjustments based on performance goals. For example, the partition management function of EWLM adjusts processor resources based on workload performance goals. Thus, EWLM can adjust the processing capacity for AIX and Linux logical partitions.

**Limitations:**

- Do not add the management partition to the partition workload group. To manage logical partition resources, workload management tools often require that you install some type of management or agent software on the logical partitions. To avoid creating an unsupported environment, do not install additional software on the management partition.
- For AIX and Linux partitions, the DLPAR support of the operating system is not the same as the DLPAR capabilities that are in the partition properties for a logical partition. The DLPAR support of the operating system reflects what each operating system supports with regard to DLPAR functions. AIX and Linux support DLPAR of processors, memory, and I/O. The DLPAR capabilities that are shown in the partition properties for a logical partition reflect a combination of the following:

  - A Resource Monitoring and Control (RMC) connection between the management partition and the client logical partition
  - The operating system's support of DLPAR

  For example, an AIX client logical partition does not have an RMC connection to the management partition, but AIX supports DLPAR of processors, memory, and I/O. In this situation, the DLPAR capabilities shown in the partition properties for the AIX logical partition indicate that the AIX logical partition is not capable of processor, memory, or I/O DLPAR. However, because AIX supports DLPAR of processors, memory, and I/O, a workload management tool can dynamically manage its resources. Workload management tools are not dependent on RMC connections to dynamically manage logical partition resources.
- If a logical partition is part of the partition workload group, you cannot dynamically manage its resources from the Integrated Virtualization Manager because the workload management tool is in control of dynamic resource management. Not all workload management tools dynamically manage processor, memory, and I/O resources. When you implement a workload management tool that manages only one resource type, you limit your ability to dynamically manage the other resource types. For example, EWLM dynamically manages processor resources, but not memory or I/O. AIX supports processor, memory, and I/O DLPAR. EWLM controls dynamic resource management of processor resources, memory, and I/O for the AIX logical partition, but EWLM does not dynamically manage memory or I/O. Because EWLM has control of dynamic resource management, you cannot dynamically manage memory or I/O for the AIX logical partition from the Integrated Virtualization Manager.

To add a logical partition to the partition workload group, complete the following steps:

1. From the Partition Management menu, click **View/Modify Partitions**. The View/Modify Partitions window displays.
2. Select the logical partition that you want to include in the partition workload group.
3. From the Tasks menu, select **Properties**. The **Partition Properties** window is displayed.
4. In the General tab, select **Partition workload group participant** and click **OK**.

**Related information**

Enabling partition management

## Reference codes for Partition Mobility

Partition Mobility reference codes are diagnostic aids that help you determine the source of a hardware or operating system problem. Using reference codes, you can find the correct solution to fix the problem. To use reference codes effectively, you must use them in conjunction with other service and support procedures.

To help you better understand your problem and how you can fix it, refer to the System Reference Codes PDF for your server model.

# Chapter 6. Partition Load Manager for AIX

The Partition Load Manager for AIX 5L provides automated processor and memory resource management across logical partitions that are capable of dynamic logical partitioning on AIX 5L.

The Partition Load Manager allocates resources to partitions on demand within the constraints of a user-defined policy. Partitions with a high demand for resources are given resources from partitions with a lower demand, improving the overall resource utilization of the system. Resources that would otherwise be unused, if left allocated to a partition that was not using them, can now be used to meet resource demands of other partitions in the same system.

The Partition Load Manager uses a client/server model to report and manage resource utilization. The clients, or *managed partitions*, notify the Partition Load Manager server when resources are either not used enough or are overused. Upon notification of one of these events, the Partition Load Manager server makes resource allocation decisions based on a user-defined resource management policy. This policy determines how much of the available resources are to be allocated to each partition.

The Partition Load Manager works much like any other system management software in that you can use it to view the resources across your partitions, group those resources into manageable segments, and allocate and reallocate those resources within or across the groups. It also locally logs activity on the partitions. The underlying processes of the Partition Load Manager rely on Resource Monitoring and Control (RMC) for network communication with the managed partitions.

Requirements for using the Partition Load Manager server include the following:

- A Hardware Management Console (HMC) must be attached to the managed system.
- The Partition Load Manager system can be running AIX 5L Version 5.2 with the 5200-04 Technology Level or AIX 5L Version 5.3.
- The Partition Load Manager system can be a partition in the managed server, a partition in a different server, or a stand-alone AIX system.
- The Partition Load Manager server system requires network connectivity to the HMC and to every managed partition.
- Multiple Partition Load Manager servers might be run on one AIX system.
- One Partition Load Manager server can manage partitions within only one managed server.

## Partition Load Manager for AIX

The Partition Load Manager provides processor and memory resource management and monitoring across logical partitions within a single managed system that uses POWER5 technology.

Partition Load Manager allows you to more effectively use resources by allowing you to set thresholds for designated resources. When a threshold is exceeded, Partition Load Manager can try to assign resources to that logical partition by using resources assigned to other logical partitions that are not being used.

Partition Load Manager is available as part of the PowerVM Editions feature.

When the Partition Load Manager resource manager starts, it registers several events on every required logical partition’s node. The following events are registered individually on all of the managed logical partitions nodes:

- Memory-page-steal high thresholds
- Memory-usage high thresholds and low thresholds

- Processor-load-average high thresholds and low thresholds

The Partition Load Manager resource manager tracks these threshold values. Every time a threshold is exceeded, Partition Load Manager receives a Resource Monitoring and Control (RMC) event. When a high threshold is exceeded, the node needs more resources. Alternately, when a low threshold is crossed, the node has more resources available than it is currently using.

When a node requests additional resources, Partition Load Manager determines whether the node can accept additional resources. If the node can accept additional resources, Partition Load Manager conducts a search for available resources. Such additional resources can be found in the following places:

- The *free pool*, which is the list of currently unused resources maintained by Partition Load Manager. These resources are reallocated freely.
- Nodes that have indicated through events that they can release resources. These resources are removed from the node that does not require them and reallocated to the node that is requesting additional resources.
- Taken away from a node that has a lesser need for the resource, or a lower priority, than the node requesting the resource. These resources are removed from the node that has lower priority and reallocated to the node that is requesting additional resources.

Determining which node is more or less deserving of resources is primarily done by taking into account certain values defined in a *policy file*. This policy file details partitions, their entitlements, their thresholds, and organizes the partitions into groups. Every node, but not every logical partition, managed by Partition Load Manager must be defined in the policy file, along with several associated attribute values. Some of the attributes that are associated with the node are the maximum, minimum, and guaranteed resource values, variable share values, and so on. Partition Load Manager takes these attributes into account when a decision is made as to whether a resource is reallocated from one logical partition to another.

For example, a machine is likely to lose its resource to a node with a higher variable shares attribute value if that machine has a lesser variable shares attribute value and currently has more resource than the guaranteed resource value given in the policy file.

## Preparing to install the Partition Load Manager

Use this procedure to prepare to install the Partition Load Manager.

Before you install the Partition Load Manager, complete the following steps:

**Name resolution**

Resolve the host name by completing the following steps:

1. Set the host name on each logical partition to the fully qualified host name, such as lpar1.domain.com.
2. If you are not using a name server, edit the **/etc/hosts** file on each logical partition to include the Partition Load Manager server host name, similar to the following:
```
172.16.0.30     lpar1.domain.com          lpar1
172.16.0.100    plmserver1.domain.com     plmserver1
```
3. If you are not using a name server, edit the **/etc/hosts** file on the Partition Load Manager server to include the logical partitions and HMC host names, similar to the following:
```
172.16.0.100    plmserver1.domain.com     plmserver1
172.16.0.30     lpar1.domain.com          lpar1
172.16.0.33     lpar2.domain.com          lpar2
172.16.0.3      p5hmc1.domain.com         p5hmc1
```

**Dynamic logical partitioning capability on logical partitions**

Determine the dynamic logical partitioning capability of logical partitions by completing the following steps:

1. To determine whether each logical partition is capable of dynamic logical partitioning, run the following command:

   ```
   lssrc -a | grep rsct
   ```

   If the Partition Load Manager resource manager daemon is running, then the logical partition has an active Resource Monitoring and Control (RMC) session with the HMC and is capable of dynamic logical partitioning.

   If the Partition Load Manager resource manager daemon is not running, check the name resolution and the network connectivity between the HMC and the LPAR.
2. If you changed the host name without rebooting, recycle the RMC daemons on each logical partition by running the following commands:

   ```
   /usr/sbin/rcst/bin/rmcctrl -z
   /usr/sbin/rsct/bin/rmcctrl -s
   ```

### RSH and RCP access to managed logical partitions from the Partition Load Manager server

Remote shell (rsh) and remote control panel (rcp) access is required to all logical partitions for setting up the Partition Load Manager. If rsh and rcp have been disabled for security reasons, use the following steps to enable these services:

1. Edit the **.rhosts** file on each logical partition to add the following lines:

   ```
   plmserver1 root
   plmserver1.domain.com root
   ```
2. Enable rsh and rcp on each logical partition by running the following commands:

   ```
   chmod 4554 /usr/sbin/rshd
   chmod 4554 /usr/bin/rcp
   ```
3. Edit the **/etc/inetd.conf** file, and uncomment the following line:

   ```
   shell stream tcp6 nowait root /usr/sbin/rshd rshd
   ```
4. Restart the inetd daemon by running the following command.

   ```
   refresh -s inetd
   ```
5. Test the rsh access from the Partition Load Manager server to each logical partition by running the following commands:

   ```
   rsh lpar1 -l root date
   rsh lpar2 -l root date
   ```

### Create an AIX user ID for the Partition Load Manager

The Partition Load Manager server is a *setuid* program that runs under the configured user ID. This user must exchange ssh keys with the configured HMC user and be authorized with Resource Monitoring and Control (RMC) before running Partition Load Manager. Use any of the management interfaces to create the **plmuser** ID on the Partition Load Manager server.

## Installing OpenSSH software tools

Use this procedure to download and install OpenSSH software tools on an AIX logical partition. OpenSSH must be set up so that you can facilitate authentication and communication between the Partition Load Manager server and the controlling Hardware Management Console (HMC).

Whenever the Partition Load Manager satisfies a resource request, it uses remote HMC commands to gather partition information and initiate dynamic logical partitioning operations. The HMC must be enabled for OpenSSH by activating the Enable/Disable Remote Command Execution task on the HMC.

When you are setting up a user on the HMC for OpenSSH, specify one of the following roles:

- System administrator
- Service representative
- Advanced operator

Before you can use OpenSSH, there must be a user on the HMC that has remote command enabled. This user must exchange ssh keys with the configured HMC user, but does not have to be the same user as the **plmuser** ID.

OpenSSH software tools support the SSH1 and SSH2 protocols. The tools provide shell functions where network traffic is encrypted and authenticated. OpenSSH is based on client and server architecture. OpenSSH runs the sshd daemon process on the AIX host and waits for the connection from clients. It supports public-key and private-key pairs for authentication and encryption of channels to ensure secure network connections and host-based authentication. For more information about OpenSSH, including the man pages, see http://www.openssh.org.

The OpenSSH software is included on the AIX 5.3 Expansion Pack. This version of OpenSSH is compiled and packaged as **installp** packages using the **openssh-3.7.1p2** level of source code. The **installp** packages include the man pages and the translated message filesets. The OpenSSH program contained in the Expansion Pack CD-ROM media is licensed under the terms and conditions of the International Program License Agreement (IPLA) for Non-Warranted Programs.

Before installing the OpenSSH **installp** format packages, you must install the Open Secure Sockets Layer (OpenSSL) software that contains the encrypted library.

After you download the OpenSSL package, you can install OpenSSL and OpenSSH.

1. Install the OpenSSL RPM package using the geninstall command, as follows:

   ```
   # geninstall -d/directory R:openssl-0.9.6g
   ```

   where *directory* is the name of the directory to which you downloaded the OpenSSL package. Output similar to the following displays:

   ```
   SUCCESSES
   ---------
   openssl-0.9.6g-3
   ```

2. Install the OpenSSH **installp** packages using the geninstall command, as follows:

   ```
   # geninstall -Y -d/directory I:openssh.base
   ```

   Use the **-Y** flag to accept the OpenSSH license agreement after you have reviewed the license agreement.

   To view the license agreement, type the following command:

   ```
   # geninstall -IapE -ddirectory openssh.base 2>&1 |pg
   ```

   After you accept the license agreement, output similar to the following displays:

   ```
   Installation Summary
   --------------------
   Name                        Level           Part        Event       Result
   -------------------------------------------------------------------------------
   openssh.base.client         3.6.0.5200      USR         APPLY       SUCCESS
   openssh.base.server         3.6.0.5200      USR         APPLY       SUCCESS
   openssh.base.client         3.6.0.5200      ROOT        APPLY       SUCCESS
   openssh.base.server         3.6.0.5200      ROOT        APPLY       SUCCESS
   ```

You can also use the **smitty license_on_media** fast path to view the license, and the **smitty install_software** fast path to install OpenSSL and OpenSSH.

The following OpenSSH binary files are installed as a result of the preceding procedure:

**scp** A file copy program similar to **rcp**

**sftp** A program similar to **FTP** that works over the SSH1 and SSH2 protocol

**sftp-server**
A SFTP server subsystem (started automatically by **sshd** daemon)

**ssh** Similar to the **rlogin** and **rsh** client programs

**ssh-add**
A tool that adds keys to **ssh-agent**

**ssh-agent**
An agent that can store private keys

**ssh-keygen**
A key-generation tool

**ssh-keyscan**
A utility for gathering public host keys from a number of hosts

**ssh-keysign**
A utility for host-based authentication

**sshd** A daemon that permits you to log in

### SSH access to the HMC from the Partition Load Manager server

After you have installed SSH, you can generate the SSH keys and communicate with the HMC.

If you are going to run the Partition Load Manager server under the **plmuser** ID, grant SSH access to the HMC from the Partition Load Manager server by using the following steps:

1. Log in under the **plmuser** ID.
2. Generate SSH keys on the Partition Load Manager server by using the following command:
   ```
   ssh-keygen -t rsa
   ```
3. Exchange SSH keys with the HMC by using the following commands:
   ```
   scp hscroot@p5hmc1:.ssh/authorized_keys2 ~/.ssh/tmp_authorized_keys2
   cat ~/.ssh/id_rsa.pub >> ~/.ssh/tmp_authorized_keys2
   scp ~/.ssh/tmp_authorized_keys2 hscroot@p5hmc1:.ssh/authorized_keys2
   ```
4. Test the SSH access to the HMC as the **plmuser** ID without using a password by using the following command:
   ```
   ssh hscroot@p5hmc1 date
   ```
5. Obtain the name of the managed system from the HMC by using the following command:
   ```
   ssh hscroot@p5hmc1 lssyscfg -r sys
   ```

   Unless the name of the managed system is changed on the HMC using the **Properties** tab on the managed system, the default managed system name is similar to the following:

   server-9117-570-SNxxxxxxx

   **Note:** The HMC hostname used in the setup and the managed system name are used in the Partition Load Manager policy. If there is more than one managed system, determine which system contains the partitions to be managed. For each managed system, use the following command:
   ```
   ssh hmcuser@hmchost lssyscfg –r lpar –m machine
   ```

   **Related information**

   AIX Toolbox for Linux Applications

## Installing the Partition Load Manager server

Use this procedure to install the Partition Load Manager server on an AIX logical partition.

To install the Partition Load Manager server, complete the following steps:

1. Mount the *Partition Load Manager* CD to your system.
2. Using either the installp command or the **smitty install_latest** fastpath, install the following filesets:
   - plm.license
   - plm.server.rte
   - plm.sysmgt.websm
   - plm.msg.en_US.server
   - plm.msg.en_US.websm
3. Read and accept the license.

Now that the Partition Load Manager server is installed, you can create a policy file and configure Resource Monitoring and Control (RMC) for the Partition Load Manager. If you create the policy file first and Web-based System Manager is being used, you can use the policy file to input the list of partitions being managed.

## Configuring the policy file

Use this procedure to configure the policy file for the Partition Load Manager server.

**Policy file concepts**

The system uses the policy file to determine which processor and memory resources may be managed by the Partition Load Manager server. The policy also includes resource shares, group definitions, and tunable parameters. This file defines the partitions that are to be managed, their guaranteed entitlements, and their minimum and maximum entitlements.

The policy file is divided into stanzas. Each stanza has a **type** field. Every stanza follows the following format:

```
<stanza_label>:
            attribute=<value>
            attribute2=<value>
            type=<value>
```

The policy file has the following rules:

- The policy file consists of a number of stanzas containing attributes.
- Stanza names may not contain any blanks and must be followed immediately by a colon (:). Only white space or a comment can follow the stanza name. For improved readability, enter stanza names starting in column 1 on the line. The following are the supported stanza types:
  - globals:
  - tunables:
  - group_name:
  - partition_name:
- Attributes consist of a name and a value separated by an equal sign (=). Attribute names and values may not contain any blanks. Only white space or a comment may follow the value. For improved readability, enter attributes so that they are indented under the containing stanza name.
- Do not repeat attributes in a stanza. Only the first attribute in a stanza is used.
- Comments begin with a number sign (#). Comments can be started in any column on the line and continue until end of line.

- Stanzas may be placed in the policy file in any order. The following is a suggested order:
  1. globals stanza
  2. tunables stanza
  3. group stanza for first group
  4. partition stanzas for partitions in first group
  5. repeat group/partition stanza for subsequent groups

The available types of stanzas and their attributes are described as follows:

**globals stanza:**

This stanza specifies global environment attributes for the Partition Load Manager server. Only one globals stanza can be specified in a Partition Load Manager policy.

The following attributes are required in the globals stanza:

| Attribute | Description |
|---|---|
| **hmc_host_name** | Host name of the Hardware Management Console (HMC) that manages the server that contains the managed partitions. This is the host name that was used for the HMC when exchanging ssh keys. |
| **hmc_cec_name** | The HMC managed system name for the server that contains the managed partitions. |
| **hmc_user_name** | The user name that the Partition Load Manager uses to send OpenSSH commands to the HMC |

The following attribute is optional in the globals stanza:

| Attribute | Minimum value | Maximum value | Default value | Description |
|---|---|---|---|---|
| **hmc_command_wait** | 1 minute | 60 minutes | 5 minutes | The number of minutes that the Partition Load Manager waits before timing out an HMC command. This is the DR Phase Timeout, one of three phases. |

**tunables stanza:**

This optional stanza is used to specify tunable attributes for the managed partitions. There are no required attributes in the tunables stanza. The Partition Load Manager has selected default values for these attributes that are appropriate for most installations. However, installations with special requirements can customize their installation by specifying the attributes in this stanza. The attributes in the tunables stanza can also be specified in the group and partition stanzas. A tunable attribute for a partition is obtained in the following order:

1. From the partition stanza.
2. From the group stanza containing the partition if tunable attribute is not specified in the partition stanza.
3. From the tunables stanza if tunable attribute is not specified in the partition or group stanzas.
4. Default value is used if tunable attribute is not specified in the partition, group, or tunables stanzas.

Specify any of the following processor-related attributes:

| Attribute | Minimum value | Maximum value | Default value | Description |
|---|---|---|---|---|
| **cpu_intervals** | 1 | 100 | 6 | The number of 10- second periods that a CPU-related sample must cross before the Partition Load Manager will activate. Setting this value higher causes the Partition Load Manager to react more slowly to system changes. Setting it lower causes the Partition Load Manager to activate more quickly. |
| **cpu_load_low** | 0.10 | 1.00 | 0.5 | The CPU load average low threshold value. A partition with a load average below this value is considered to have unneeded CPU capacity.<br>**Note:** The minimum delta between **cpu_load_low** and **cpu_load_high** is 0.10. |
| **cpu_load_high** | 0.2 | 10.0 | 1.0 | The CPU load average high threshold value. A partition with a load average above this value is considered to need more CPU capacity.<br>**Note:** The minimum delta between **cpu_load_low** and **cpu_load_high** is 0.10. |
| **cpu_free_unused** | | | No | Indicates whether CPU capacity not needed by a partition is removed from the partition. A value of no indicates unneeded CPU capacity remains in the partition until another partition has a need for it. A value of yes indicates unneeded CPU capacity is removed from the partition when the partition no longer has a need for it. |

Specify any of the following shared processor-related attributes:

| Attribute | Minimum value | Maximum value | Default value | Description |
|---|---|---|---|---|
| **ec_delta** | 1 | 100 | 10 | The amount of CPU entitled capacity to add or remove from a shared processor partition. The value specifies the percentage of the partition's current entitled capacity to add or remove. |
| **ec_per_vp_min** | 0.1 | 0.9 | 0.5 | The minimum amount of entitled capacity per virtual processor. This attribute prevents a partition from having degraded performance by having too many virtual processors relative to its entitled capacity. When entitled capacity is removed from a partition, virtual processors will also be removed if the amount of entitled capacity for each virtual processor falls below this number.<br>**Note:** The minimum delta between **ec_per_vp_min** and **ec_per_vp_max** is 0.10. |

| Attribute | Minimum value | Maximum value | Default value | Description |
|---|---|---|---|---|
| **ec_per_vp_max** | 0.2 | 1.0 | 0.8 | The maximum amount of entitled capacity per virtual processor. This attribute controls the amount of available capacity that may be used by an uncapped shared CPU partition. When entitled capacity is added to a partition, virtual processors will be added if the amount of the entitled capacity for each virtual processor exceeds this number. Increasing the number of virtual processors in an uncapped partition allows the partition to use more of the available CPU capacity.<br>**Note:** The minimum delta between **ec_per_vp_min** and **ec_per_vp_max** is 0.10. |

Specify any of the following memory-related attributes:

| Attribute | Minimum value | Maximum value | Default value | Description |
|---|---|---|---|---|
| **mem_intervals** | 1 | 100 | 6 | The number of 10-second periods that a memory-related sample must cross before the Partition Load Manager will activate. Setting this value higher causes the Partition Load Manager to react more slowly to system changes. Setting it lower causes the Partition Load Manager to activate more quickly. |
| **mem_util_low** | 1 | 90 | 50 | The memory utilization low threshold value. A partition with a memory utilization below this value is considered to have unneeded memory. Units are expressed as a percent.<br>**Note:** The minimum delta between **mem_util_low** and **mem_util_high** is 10. |
| **mem_util_high** | 1 | 100 | 90 | The memory utilization high threshold value. A partition with a memory utilization above this value is considered to need more memory. Units are expressed as a percent.<br>**Note:** The minimum delta between **mem_util_low** and **mem_util_high** is 10. |
| **mem_pgstl_high** | 0 | 2147483647 | 0 | The page steal threshold. A partition with a page steal rate, which is the number of page steals per second, greater than or equal to this value is considered to need more memory. Units are expressed as an integer value. The result of checking this threshold is logically ANDed with the result of the **mem_util_high** threshold check when determining if memory is needed. |

| Attribute | Minimum value | Maximum value | Default value | Description |
|---|---|---|---|---|
| **mem_free_unused** | | | No | Indicates when memory not needed by a partition is removed from the partition. A value of no indicates unneeded memory remains in the partition until another partition has a need for it. A Yes value indicates unneeded memory is removed from a partition when the partition no longer has a need for it. |
| **mem_delta** | 1 | 256 | Specifies one LMB to be removed or added to a partition at a time | The amount of memory to be removed or added to a partition. The units are in megabytes. If the value is less than the system's logical memory block (LMB) size, the value is rounded up to the system's LMB size. If the value is greater than the system's LMB size but not a multiple of LMB size, the value is rounded down to the nearest LMB multiple size. |

**group_name stanza:**

This stanza specifies the name and global attributes for a group, and any or all of the tunables stanzas. The name on a group stanza specifies the name of the group. The group stanza allows you to create multiple groups of partitions that are managed independently. At least one group must be defined.

The following attributes are required in the group stanza:

- type = group
- cpu_maximum
- mem_maximum

The **cpu_maximum** attribute specifies if processor management is desired for the partitions in the group and if desired the amount of processor capacity to be allocated to the partitions. If processor management is specified, processor management is performed for all partitions in the group. Specifying a **cpu_maximum** value of 0 specifies processor management is not performed for the partitions in the group.

All partitions in a group must have the same processor type. The **cpu_type** attribute specifies the processor type for all the partitions in the group and is written as follows:

```
cpu_type = dedicated | shared
```

The **mem_maximum** attribute specifies memory management is desired for the partitions in the group and if desired the amount of memory to be allocated to the partitions. If memory management is specified, memory management is performed for all partitions in the group. Specifying a **mem_maximum** value of 0 specifies memory management is not performed for the partitions in the group.

You can specify **cpu_maximum** and **mem_maximum** values greater than the amount of physical resources in the server. In this situation, all available resources will be used to satisfy resource requests for the managed partitions.

The following attributes are required in this stanza:

| Attribute | Description |
|---|---|
| **type=group** | An attribute identifying this as a group stanza. The attribute must be specified as `type = group`. |

| Attribute | Description |
|---|---|
| **cpu_maximum** | The maximum amount of CPU capacity to be allocated to partitions in the group. The units are in physical CPU units. A value of 0 indicates CPUs are not managed for the partitions in the group. |
| **mem_maximum** | The maximum amount of memory to be allocated to partitions in the group. The units are in megabytes (MB). A value of 0 indicates memory is not be managed for the partitions in the group. |
| **cpu_type** | The processor type of the partitions in the group. All partitions in the group must be the same type. The attribute value must either be dedicated or shared. |

**partition_name stanza:**

This stanza specifies the name and attributes for a partition. A partition stanza is required for every managed partition.

The name of the partition stanza is the host name of the managed partition.

The following attributes are required in a partition stanza:

- type = partition
- group = group_name

The following attributes are optional in the partition stanza:

- cpu_minimum
- cpu_guaranteed
- cpu_maximum
- cpu_shares
- mem_minimum
- mem_guaranteed
- mem_maximum
- mem_shares

If not specified, the **cpu_minimum**, **cpu_guaranteed**, and **cpu_maximum** attribute values are obtained from the CPU minimum, desired, and maximum HMC partition definition values respectively. Similarly, the **mem_minimum**, **mem_guaranteed**, and **mem_maximum** attribute values are obtained from the minimum, desired, and maximum HMC partition memory definition values. The shares values default to 1.

If minimum, guaranteed, and maximum values are specified in the policy, the values must satisfy the following relationship:

minimum <= guaranteed <= maximum

If management of CPU or memory resource is not wanted in a specific partition in a group, the values for the resource can all be specified as the same value. If management of CPU or memory resource is not wanted for all partitions in a group, the **cpu_maximum** or **mem_maximum** attributes in the group definition can be set to 0.

Any CPU or memory values specified in the policy must be compatible with the partition's HMC partition definition. You cannot use the Partition Load Manager to decrease a partition's minimum below the HMC minimum. Nor can you use the Partition Load Manager to increase a partition's maximum over the HMC maximum. System administrators are responsible for ensuring that the Partition Load Manager policies and HMC partition definitions are compatible.

The **cpu_shares** and **mem_shares** attributes are optional in the partition stanza, with default values set to 1.

The default value for **cpu_shares** is to have equal shares for all partitions in the group. The default **cpu_shares** value for shared, uncapped processor partitions is not obtained from the variable weight attribute of the partition's HMC definition. If the **cpu_shares** attribute is not specified, the Partition Load Manager does not set the variable weight HMC attribute for the partition. (The variable weight value set by the HMC continues to be used.) If the **cpu_shares** attribute is specified and the partition is shared or uncapped, the Partition Load Manager sets the partition's variable weight HMC attribute to the **cpu_shares** value.

The following tunable attributes are used in the partition stanza:

| Attribute | Minimum value | Maximum value | Default value | Description |
|---|---|---|---|---|
| **type** | | | | A required attribute identifying this as a partition stanza. The attribute must be specified as `type = partition`. |
| **group** | | | | A required attribute specifying the group containing this partition. |
| **cpu_minimum** | | | | The minimum amount of CPU capacity to be allocated to a partition. The units are in physical CPU units. |
| **cpu_guaranteed** | | | | The guaranteed amount of CPU capacity to be allocated to a partition. The units are in physical CPU units. |
| **cpu_maximum** | | | | The maximum amount of CPU capacity to be allocated to partition. The units are in physical CPU units. |
| **cpu_shares** | 0 | 255 | 1 | A factor without units that is used to specify how available CPU capacity in excess of the **cpu_guaranteed** is distributed to partitions in the group. The available excess CPU capacity is allocated to partitions using the following formula:<br>(**cpu_shares**) / (sum of **cpu_shares** from active partitions in the group)<br>**Note:** Specifying a minimum value of 0 limits a partition to receiving only its **cpu_guaranteed** amount of CPU capacity. |
| **mem_minimum** | | | | The minimum amount of memory to be allocated to the partition. The units are in megabytes (MB). |
| **mem_guaranteed** | | | | The guaranteed amount of memory to be allocated to the partition. The units are in megabytes (MB). |
| **mem_maximum** | | | | The maximum amount of memory to be allocated to the partition. The units are in megabytes (MB). |

| Attribute | Minimum value | Maximum value | Default value | Description |
|---|---|---|---|---|
| **mem_shares** | 0 | 255 | 1 | A factor with no units that is used to specify how available memory in excess of the **mem_guaranteed** is distributed to the partitions in the group. The available excess memory is allocated to partitions using the following formula:<br>(mem_shares) /<br>(sum of mem_shares<br>from competing partitions)<br>**Note:** Specifying a minimum value of 0 limits a partition to receiving only its **mem_guaranteed** amount of memory. |

**Example of policy file creation and configuration**

Using the Web-based System Manager, create a policy file by using the following steps as an example.

**Note:** If you are using a remote X server, set the *DISPLAY* variable, and use the wsm & command to start the Web-based System Manager client.

1. Create a policy file.
2. Add the policy file name: **/etc/plm/policies/plm1**
3. Add the following global values for the following fields:

   Hardware Management Console (HMC) name: **p5hmc1**

   HMC user name: **hscroot**

   Central Electronic Complex name: **server-9117-570-SNxxxxxxx**
4. Obtain the names of the LPARs and settings from the HMC by running the following commands:
   - `ssh hscroot@p5hmc1 lssyscfg -r lpar -m server-9117-570-SNxxxxxxx` (LPAR names and default profile names)
   - `ssh hscroot@p5hmc1 lshwres -r proc -m server-9117-570-SNxxxxxxx --level lpar` (settings)
   - `ssh hscroot@p5hmc1 lshwres -r proc -m server-9117-570-SNxxxxxxx --level sys` (system resources)

   The output includes the following information:
   - `name=lpar1, default_profile=default`
   - `curr_min_proc_units=0.5, curr_proc_units=0.75, curr_max_proc_units=1.25`
   - `name=lpar2, default_profile=default`
   - `curr_min_proc_units=0.5, curr_proc_units=0.75, curr_max_proc_units=1.25`
5. Add the following group information to the policy file:

   Group name: **plm1**

   Maximum CPU: 1.75

   Maximum Memory: N/A

   CPU type: shared

   Select **CPU management**

   Deselect **Memory management**
6. Add the following information for partitions for CPU resource management:
   - Partition name: lpar1.domain.com (this is the fully qualified host name for lpar1)

- Group name: **plm1**
- Resource Entitlements:

  Minimum CPU: 0.5<br>
  Guaranteed CPU: 0.75<br>
  Maximum CPU: 1.25<br>
  CPU variable shares: 1 (default)
- Partition name: lpar2.domain.com
- Group name: **plm1**
- Resource Entitlements:

  Minimum CPU: 0.5<br>
  Guaranteed CPU: 0.75<br>
  Maximum CPU: 1.25<br>
  CPU variable shares: 1 (default)
- Tunable attributes:

  CPU load average high threshold: 0.8<br>
  CPU load average low threshold: 0.2

## Querying partition status

You can use Partition Load Manager to query the status of the logical partitions on your managed system.

Any user can run the **xlplm** command to obtain status information for running instances of Partition Load Manager.

**Query the status of Partition Load Manager**

To query the status of all running instances of Partition Load Manager, type the following command:

```
xlplm -Q
```

A list of the instances that are running is displayed. If there are no instances running, no output is displayed.

**Query the attributes of an instance**

To query the attributes of a single instance, type the following command, where *test1* is the name of the instance:

```
xlplm -Q test1
```

Output from this command will be similar to the following:

```
PLM Instance: test1

GROUP: group1
          CUR       MAX     AVAIL     RESVD      MNGD
CPU:      6.00      4.00      0.00      0.00       Yes
MEM:      8192      8192         0         0       Yes

thimblelp10.server.company.com

RESOURCES:
          CUR       MIN      GUAR       MAX       SHR
CPU:      3.00      1.00      3.00      3.00         1
MEM:      4096      1024      4096      4096         1
```

```
thimblelp11.server.company.com

RESOURCES:
          CUR       MIN      GUAR       MAX       SHR
CPU:      3.00      1.00      3.00      3.00         1
MEM:      4096      1024      4096      4096         1
```

### View additional information from a query

To view additional information from the query of a specific instance, type the following command, where *test1* is the name of the instance:

```
xlplm -v -Q test1
```

The verbose output from this command will be similar to the following:

```
PLM Instance: test1

 CEC Name            Thimble
 Mode                monitor
 Policy              /etc/plm/policies/policy1
 Log                 /tmp/log.test
 HMC Host            kbuphsc2.server.company.com
 HMC User            hscroot

GROUP: group1
          CUR       MAX     AVAIL     RESVD      MNGD
CPU:      6.00      4.00      0.00      0.00       Yes
MEM:      8192      8192         0         0       Yes

CPU TYPE: dedicated

thimblelp10.server.company.com

  RESOURCES:
          CUR       MIN      GUAR       MAX       SHR
CPU:      3.00      1.00      3.00      3.00         1
MEM:      4096      1024      4096      4096         1

  TUNABLES:
             INTVL    FRUNSD    LOADLO    LOADHI     DELTA     PGSTL
CPU:         6         0      0.40        1.00    1.00         -
MEM:         6         0       50%       90%       256         0\

thimblelp11.server.company.com

  RESOURCES:
          CUR       MIN      GUAR       MAX       SHR
CPU:      3.00      1.00      3.00      3.00         1
MEM:      4096      1024      4096      4096         1

TUNABLES:
             INTVL    FRUNSD    LOADLO    LOADHI     DELTA     PGSTL
CPU:         6         0      0.40      1.00      1.00         -
MEM:         6         0       50%       90%       256         0

*****************************************************************************
```

### Allocate resources to partitions

You can allocate resources to specific partitions and even reserve resources for specific partitions regardless of when those partitions will use the resources. You can reserve and allocate resources from a group of managed partitions using the **xlplm -R** command. Those resources that are reserved can be used to create a new unmanaged partition, or to make room for a new partition to enter the managed group.

Reserved resources will not be allocated to any existing partition in a group unless they are first released. If a previously offline partition comes online and enters a managed group, any reserved resources within that group automatically are removed from the collection of reserved resources, called the *free pool*, and assigned to the new partition. If the reserved resources are used instead to create a new, unmanaged partition, they can be released to the group after the new partition has booted and can then be automatically reclaimed by the managed group if they later become available and are needed.

The requested reservation amount is absolute, so a reserve command can result in either a reserve or a release, depending on the current reservation amount. The minimum allowed changes in the reservation amounts are the following:

- 1 MB for memory
- 1 processor unit for a dedicated processor group
- 0.01 processor unit for a shared processor group

When you reserve resources, the free pool for the target group is first checked for available resources. If the free pool has enough resources to satisfy the request, the requested amount is removed from the free pool. If the free pool does not have enough resources to satisfy the request, resources will be taken from one or more partitions with the lowest workload, or least need for the resources. A reservation request will fail if the requested amount is more than the minimum allowed for the group.

**Manage memory resource requests**

The following is an example of how to use Partition Load Manager to manage memory resource requests. This example shows how Partition Load Manager responds to memory resource requests between two partitions:

The two partitions, LP0 and LP1, are configured as follows:

```
LP0:      Minimum = 1024 MB
          Guaranteed = 1024 MB
          Maximum = 4096 MB
          Weight = 2
          Current Entitlement = 1024 MB

LP1:      Minimum = 1024 MB
          Guaranteed = 1024 MB
          Maximum = 4096 MB
          Current Entitlement = 1024 MB
          Weight = 1
```

The total amount of memory managed by Partition Load Manager is 5120 MB. With each partition's current memory allocation, shown as `Current Entitlement = 1024 MB`, Partition Load Manager assumes that the remaining 3072 MB is unallocated and available.

If both partitions become loaded in terms of memory use, then events demanding more memory resources are generated and sent to the Partition Load Manager server. For each event received, Partition Load Manager identifies the partition as a *taker*. At the same time, Partition Load Manager checks whether the partition is currently using more than its guaranteed amount. If so, the partition is identified as an *excess user*. Because there are available resources, Partition Load Manager satisfies the request immediately and allocates memory in the amount of **mem_increment** (defined either in the Partition Load Manager policy or by the internal default value) to the partition from the available memory. After the available memory is depleted, the new entitlement allocations are as follows:

```
LP0:      Current Entitlement = 2560 MB
LP1:      Current Entitlement = 2560 MB
```

Even with the current allocations, the partitions continue to generate events demanding more memory resources.

For each event, Partition Load Manager continues to identify the partition as a taker and excess user because the partition has more resources allocated than is shown as its guaranteed entitlement. However, because there are no available resources, the request is queued if there are no other resource *donors* or any other excess users. When the request from the second partition is received, it is also marked as a taker and an excess user. Because there is an excess user already queued, Partition Load Manager can satisfy the resource request.

Because both LP0 and LP1 are takers and excess users, Partition Load Manager uses the weight associated with each as the determining factor of how the extra entitlement (the sum of the current entitlement for each partition minus the sum of each partition's guaranteed allotment) will be distributed between the two partitions.

In this example, of the extra 3072 MB, the LP0 partition is allocated 2048 MB and the LP1 partition is allocated 1024 MB. Partition Load Manager assigns the **mem_incrememt** MB of memory from the LP1 partition to the LP0 partition.

With constant memory requests from each partition, Partition Load Manager eventually distributes the memory so that current entitlements become the following:

```
LP0:      Current Entitlement = 3072 MB
LP1:      Current Entitlement = 2048 MB
```

**Manage processor resources in a shared partition environment**

The following example describes how Partition Load Manager manages processor resources in a shared partition environment. The two partitions are configured as follows:

```
LP0:      Minimum = 0.1
          Guaranteed = 0.5
          Maximum = 2.0
          Max entitlement per virtual processor = 0.8
          Weight = 3
          Current entitlement = 0.1
          Current number of virtual processors = 1

LP1:      Minimum = 0.1
          Guaranteed = 0.5
          Maximum = 2.0
          Max entitlement per virtual processor = 0.8
          Weight = 1
          Current entitlement = 0.1
          Current number of virtual processors = 1
```

The total amount of processor entitlement managed by Partition Load Manager is 2.0. The amount that is currently allocated to each partition, 0.1, leaves 1.8 of unallocated processor entitlement that Partition Load Manager can distribute.

If both partitions begin running processor-intensive jobs, they request more processor entitlement by sending requests to the Partition Load Manager. Partition Load Manager then identifies the demanding partitions as takers and as excess users if the current entitlement is above its guaranteed value.

In addition to managing processor entitlement, Partition Load Manager also manages the number of virtual processors. When either partition's current entitlement exceeds 0.8, a virtual processor is also added.

In this example, Partition Load Manager assigns the available entitlement until the partitions reach the following state:

```
LP0:      Current entitlement = 1.0
          Current number of virtual processors = 2

LP1:      Current entitlement = 1.0
          Current number of virtual processors = 2
```

If the partitions continue to demand more resource, then Partition Load Manager redistributes the assigned entitlement based on the weight and excess entitlement. Here, between the LP0 partition and the LP1 partition, the total excess amount is 1.5. Because LP0 has a weight of 3 and LP1 has a weight of 1, Partition Load Manager removes processor entitlement from the LP1 partition and reassigns it to the LP0 partition. If both partitions remain busy, then the resource allocation becomes the following:

```
LP0:      Current entitlement = 1.25
          Current number of VPs = 2

LP1:      Current entitlement = 0.75
          Current number of VPs = 2
```

## Configuring Resource Monitoring and Control (RMC)

Use this procedure to configure Resource Monitoring and Control (RMC) and to verify that RMC is installed correctly.

The Partition Load Manager server uses RMC to communicate with the managed logical partitions.

The RMC setup comprises host authentication and user authorization. The host authentication involves a public key exchange between the Partition Load Manager server and the managed nodes (partitions). This allows the Partition Load Manager server to connect, or create a session, to the managed system. The user authorization involves adding an entry to the RMC ACL (Access Control) file and allows the **plmuser** (the Partition Load Manager server) access to the required resource class. The **plmsetup** script automates these tasks using remote shell commands. If the remote shell is unavailable or not configured, the administrator can perform these tasks manually.

Run the following shell script as the root user on the managing machine that will run the Partition Load Manager:

```
/etc/plm/setup/plmsetup
```

After the script runs successfully, the RMC ACL file on the remote machine will have an entry similar to the following:

```
vendor.LPAR           plmuser@plmserver1.domain.com      *       rw
```

The setup procedure takes the following as arguments:

- The user ID under which the Partition Load Manager is to run
- The host name of the partition

This user ID is used to set up the RMC ACL files on the logical partitions. ACL files are used to authenticate authorized users for each resource class when they connect to the RMC subsystem. Only this user will be permitted access to the Partition Load Manager. Only the authorized user can run Partition Load Manager. Any user is able to run commands that only display data.

**Resource Monitoring and Control (RMC) configuration for the Partition Load Manager**

Configure RMC for the Partition Load Manager by doing the following steps.

1. Select **Set up Management of Logical Partitions**.

   Authenticated user name: **plmuser**
2. Select **Automatically setup with each partition in the policy file**.

   Policy file name: **/etc/plm/policies/plm1**

3. Click **OK**.

This configuration can also be done using the command line if you are the root user on the Partition Load Manager server:

```
/etc/plm/setup/plmsetup lpar_hostname plmuser
```

To run this command, you must have rsh and rcp access. After the setup has been run, you can delete the **.rhosts** file.

## Verifying the Resource Monitoring and Control (RMC) setup

Use this procedure to verify the Resource Monitoring and Control (RMC) setup.

To verify the RMC setup, run the following as the Partition Load Manager user for each of the logical partitions that were used with the plmsetup script. Replace *PART_HOST* with the name of the logical partitions in the following command:

```
 CT_CONTACT=PART_HOST lsrsrc vendor.LPAR
```

If the persistent attributes of the resource class are displayed, then verification is successful.

If the persistent attributes of the resource class are not displayed, try the following steps:

- To troubleshoot host or connection errors, complete the following steps.
  1. Perform host-based authentication. Complete the following steps:
     a. Run the following command on both the Partition Load Manager server machine and the logical partition.

        ```
        /usr/sbin/rsct/bin/ctsvhbal
        ```

        A list of identities are displayed. These are identities as which the known partition host can be identified.

     b. Run the following command on both the Partition Load Manager server machine and the logical partition.

        ```
        /usr/sbin/rsct/bin/ctsthl -l
        ```

        On the Partition Load Manager server machine, there is an entry for the logical partition. On the logical partition, there is an entry for the Partition Load Manager server machine. The *HOST_IDENTITY* value must match one of the identities listed in the respective ctsvhbal command output.
  2. If the *HOST_IDENTITY* value in the ctsthl command output does not match the correct identity in the ctsvhbal command output on either the Partition Load Manager server machine or the logical partition , change the *HOST_IDENTITY* value by completing the following steps:
     a. Remove the incorrect *HOST_IDENTITY* value by running the following command:

        ```
        /usr/sbin/rsct/bin/ctsthl -d -n HOST_IDENTITY
        ```

     b. Add the correct *HOST_IDENTITY* value by running the following command:

        ```
        /usr/sbin/rsct/bin/ctsthl -a -n IDENTITY -m METHOD \  -p ID_VALUE
        ```

        The value for the **METHOD** parameter can be obtained from the ctsthl command. Look for an entry for the machine itself. In that entry, use the value in the Identifier Generation Method field. One example is `rsa512`. For the **ID_VALUE** parameter value, use the Identifier Value field in the same entry.
- To troubleshoot user or authorization type errors, check the ACL file on the logical partition. In the /var/ct/cfg/ctrmc.acls file, there is a stanza for `vendor.LPAR` towards the end of the file that looks similar to the following:

  ```
  vendor.LPAR          plmuser@plmserver1.domain.com      *       rw
  ```

  The user name in the stanza must match the actual user name to run the Partition Load Manager. Also, the host name in the stanza must match what was returned by the ctsvhbal command which was run on the Partition Load Manager server machine. If the host name is incorrect, run the plmsetup script again, this time using the *IDENTITY* provided by the ctsvhbal command.

For additional information about cluster configuration and management, see the Cluster library Web site.

**Related information**

Cluster library

## Starting and stopping the Partition Load Manager server

Use this procedure to start and stop the Partition Load Manager server and to check the Partition Load Manager statistics.

**Starting the Partition Load Manager server**

- Assume the following environment:

  Configuration name: `default`

  Policy file name: **/etc/plm/policies/plm1**

  Log file name: **/var/opt/plm/plm.log**

  Operation mode: `management` or `monitoring`
- Start the Partition Load Manager server by doing one of the following:
  - For management operation mode, type the following command from the command line:
    ```
    xlplm -S -p /etc/plm/policies/plm1 -l /var/opt/plm/plm.log -o M
    ```
  - For monitoring operation mode, type the following command from the command line:
    ```
    xlplm -S -p /etc/plm/policies/plm1 -l /var/opt/plm/plm.log -o N
    ```

Check the log for errors by typing the following command:

```
tail -f /var/opt/plm/plm.log
```

**Checking the Partition Load Manager statistics**

The xlpstat command is independent of the Partition Load Manager server and therefore can be run whether or not the Partition Load Manager server is running. The xlpstat command can be run any time after the RMC setup is complete.

Check the Partition Load Manager statistics by typing the following command, which checks the statistics every five seconds until you cancel the command:

```
xlpstat -p /etc/plm/policies/plm1 5
```

**Stopping the Partition Load Manager server**

Stop the Partition Load Manager server by doing one of the following steps:

- Assume that the configuration name is `default`.
- From the command line, type the following:
  ```
  xlplm -K default
  ```

## Commands for the Partition Load Manager

A description of each Partition Load Manager command is given here. This information is also available from the command line using the man command.

### xlplm command

The xlplm command starts, stops, modifies, reserves, and queries a Partition Load Manager server.

## Purpose

Start, stop, modify, reserve, and query a Partition Load Manager server.

This command is for use with the Partition Load Manager on AIX only.

## Syntax

xlplm -S -p *policy_file* -l *log_file* [ -o *operation_mode*] [ **configuration** ]

xlplm -K [ **configuration** ]

xlplm -M [ -p *policy_file* ] [ -l *log_file* ] [ -o *operation_mode* ] [**configuration** ]

xlplm -R -g *group_name* [ -c *cpu_resource_size* ] [ -m *memory_resource_size* ] [ **configuration** ]

xlplm -Q [ -r ] [ -f ] [ **configuration** ]

xlplm -C -p *policy_file*

## Description

The Partition Load Manager server xlplmd daemon performs the specified resource management operations.

## Flags

| Flag | Description |
|---|---|
| -c *cpu_resource_size* | Specifies the amount of processor resources to reserve. |
| -C | Verifies the validity of the policy file. |
| -f | By default, the query operation displays the active configuration values that might have been adjusted at run time due to conflicts with the partition profile. This option changes the output to display the values that were specified in the policy file. |
| -g *group_name* | Specifies the name of a group in the policy file. Use this flag when you are reserving or releasing resources. The resources that you want to reserve are removed from the specified group in the policy file. When you are releasing resources, they are placed in the free pool of the specified group in the policy file. |
| -K | Stop the Partition Load Manager instance. To use this flag, you must either have root authority or be logged in as the authorized **plmuser** user ID. |
| -l **log_file** | Specifies the name of the file you want to contain the Partition Load Manager activity log. |
| -M | Modify a Partition Load Manager server. To use this flag, you must either have root authority or be logged in as the authorized **plmuser** user ID. |
| -m *memory_resource_size* | Specifies the amount of memory resource to reserve. |
| -o *operation_mode* | Specifies whether the Partition Load Manager server is to operate in management mode, with a value of `M`, or monitoring mode, with a value of `N`. When the xlplm daemon starts, the default value is management mode, or `M`. |
| -p *policy_file* | Specifies the name of the Partition Load Manager policy file. |
| -Q | Query the Partition Load Manager server status. |
| | |

| Flag | Description |
|---|---|
| -R | Reserve or release resources from a partition managed by a Partition Load Manager server. To use this flag, you must either have root authority or be logged in as the authorized **plmuser** user ID. |
| -r | Use the Partition Load Manager server in raw data mode. |
| -S | Start a Partition Load Manager instance. To use this flag, you must either have root authority or be logged in as the authorized **plmuser** user ID. |
| | |

## Parameters

| Parameter | Description |
|---|---|
| **configuration** | Identifies an instance of the Partition Load Manager management. This parameter must be specified if there are multiple instances of the Partition Load Manager server on your system. If the parameter is not specified, a default value is used. |

## Exit status

This command returns the following exit values:

**0** Command completed successfully.

**1** The program encountered a nonrecoverable internal error, such as a memory allocation or system call failure.

**2** The specified log file could not be opened or created.

**3** The specified policy file could not be opened or created.

**4** A required temporary file could not be created in the /tmp directory.

**5** The specified policy is not valid.

**6** The daemon failed to start. This could be the result of an internal error or an inability to communicate with the Hardware Management Console (HMC).

**7** Command line usage error.

**8** The number specified for the reservation amount was not valid.

**9** The current user does not match the authorized user in the /etc/plm/auth/plmuser file, or the file could not be read.

**10** An instance with the requested name already exists.

**11** An instance with the requested name does not exist.

**12** The requested mode is the same as the current mode.

**13** A remote command to the HMC failed.

**14** A reservation request failed due to one of the following reasons:

- unknown group
- reservation amount is already set to requested amount
- could not reserve the requested amount
- the requested resource is not managed

### Examples

1. Start the Partition Load Manager server in management mode with a configuration name of cec1 by typing one of the following commands on the Partition Load Manager server system:

   ```
   cd /etc/xlplm/cec1
   xlplm -S -p policy -l log cec1
   ```

   or

   ```
   xlplm -S -p /etc/xlplm/cec1 -l /etc/xlplm/cec1/log -o M cec1
   ```
2. Start the Partition Load Manager server in monitoring mode by typing the following:

   ```
   xlplm -S -p policy -l log -o N cec1
   ```
3. Stop the Partition Load Manager server by typing the following:

   ```
   xlplm -K cec1
   ```
4. Load a new policy into the Partition Load Manager server by typing the following:

   ```
   xlplm -M -p evening_policy cec1
   ```
5. Start using a new log file, called newlog, for the Partition Load Manager server by typing the following:

   ```
   xlplm -M -l newlog cec1
   ```
6. Display configuration names for the active Partition Load Manager server by typing the following:

   ```
   xlplm -Q
   ```

## xlpstat command

The xlpstat command displays logical partition load statistics for a list of host names.

### Purpose

Displays logical partition load statistics for a list of host names.

This command is for use with the Partition Load Manager on AIX only.

### Syntax

xlpstat [-r] {-p **policy_file** | -f **host_list**} [**interval**] [**count**]

### Description

Display load statistics for one or more remote logical partitions. The command will contact the remote systems every number of seconds specified by the **interval** parameter for each number of intervals specified by the **count** parameter. If the **interval** parameter and the **count** parameter are omitted, the remote systems are queried once. If only the **count** parameter is omitted, the remote systems are queried every number of seconds specified by the **interval** parameter until the command is terminated by the user.

The caller of this command must be the root user or the Partition Load Manager authorized user.

For the formatted output, the output is displayed as follows, for each host listed in the input file:

```
                 CPU                      MEM
        ------------------------    ------------------
  STAT  TYP    CUR    PCT   LOAD     CUR    PCT  PGSTL  HOST

  up    D     4.00   50.15   0.65   1024  43.10      0  testlp1
  up    D     2.00   95.72   0.90   2048  97.48    250  testlp2
  up    D    10.00   98.31   1.03   5120  72.25      0  testlp3
```

```
STAT Partition status. May be "up" or "down".
TYP Partition type. May be "D" (dedicated) "S" (shared) or "U" (unknown).
      If the type is "U", the command was unable to query the partition type and
      there may be a connectivity or authentication problem.
CUR The current amount of resource allocated to the partition.
PCT Percent utilization for the resource
LOAD CPU load average
PGSTL Page steals per second
HOST Managed host name
```

The raw output is displayed as a header containing column descriptions followed by one line of data for each host:

```
#host_name:group_name:status:cpu_type:cpu_ent:cpu_util:cpu_load:mem_ent:mem_util:mem_pgstl
testlp1.mydomain.com:group1:up:dedicated:4.00:45.05:0.38:1024:75.00:0
testlp2.mydomain.com:group1:up:dedicated:2.00:87.23:0.92:2048:92.21:123
testlp3.mydomain.com:group1:up:dedicated:10.00:95.17:1.01:5120:70.30:0
```

## Flags

| Flag | Description |
|---|---|
| -r | Raw output mode. Data is printed in colon separated format, with one line per host. |
| -p **policy_file** | Retrieves the host list from the given policy file. |
| -f **host_list** | Retrieves the host list from the given plain text file. This file has one host name per line. |
|  |  |

## Exit status

This command returns the following exit values:

**1** Internal error.

**3** Could not open input file.

**5** Invalid policy file.

**7** Usage error.

**9** Not authorized.

# Chapter 7. PowerVM Lx86

Learn about PowerVM Lx86, its availability, and supported hardware.

The PowerVM Editions (formerly known as Advanced POWER Virtualization) hardware feature includes PowerVM Lx86. Lx86 is a dynamic, binary translator that allows Linux applications (compiled for Linux on Intel) to run without change, alongside local Linux on POWER applications. Lx86 makes this possible by dynamically translating x86 instructions to POWER and caching them to enhance translation performance. In addition, Lx86 maps Linux on Intel system calls to Linux on Power system calls. No modifications or recompilations of the x86 Linux applications are needed.

Lx86 creates a virtual x86 environment, within which, the Linux on Intel applications can run. Currently, a virtual Lx86 environment supports SUSE or Red Hat Linux x86 distributions. The translator and the virtual environment run strictly within the user-space. No modifications to the POWER kernel are required. Lx86 does not run the x86 kernel on the POWER machine. The Lx86 virtual environment is not a virtual machine. Instead, x86 applications are encapsulated so the operating environment appears to be Linux on x86, even though the underlying system is a Linux on POWER system.

Lx86 is included in the PowerVM Express Edition, PowerVM Standard Edition, and in the PowerVM Enterprise Edition.

For more information about Lx86, see PowerVM Lx86 for x86 Linux Applications Administration Guide.

# Appendix. Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products successfully.

The following list includes the major accessibility features:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are tactilely discernible and do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

# Notices

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

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

The manufacturer 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 the manufacturer.

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

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** THIS INFORMATION IS PROVIDED "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. The manufacturer 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 Web sites not owned by the manufacturer 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 product and use of those Web sites is at your own risk.

The manufacturer may use or distribute any of the information you supply 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 the manufacturer.

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 information and all licensed material available for it are provided by IBM® under terms of the IBM Customer Agreement, IBM International Program License Agreement, IBM License Agreement for Machine Code, or any equivalent agreement between us.

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 measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

All statements regarding the manufacturer's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The manufacturer's prices shown are the manufacturer'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 the names and addresses used by an actual business enterprise 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 the manufacturer, 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. The manufacturer, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

CODE LICENSE AND DISCLAIMER INFORMATION:

The manufacturer grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, THE MANUFACTURER, ITS PROGRAM DEVELOPERS AND SUPPLIERS, MAKE NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

UNDER NO CIRCUMSTANCES IS THE MANUFACTURER, ITS PROGRAM DEVELOPERS OR SUPPLIERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY:

1. LOSS OF, OR DAMAGE TO, DATA;
2. SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR
3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.

SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF DIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, SO SOME OR ALL OF THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU.

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_. All rights reserved.

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

## Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX
AIX 5L
BladeCenter

IBM

Micro-Partitioning
OpenPower
POWER
Power Architecture
POWER5
POWER6
PowerVM Editions

Tivoli
Tivoli Enterprise

Intel, Intel Inside®, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium® are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Red Hat, the Red Hat ″Shadow Man″ logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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.

## Terms and conditions

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

**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 the manufacturer.

**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 the manufacturer.

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

The manufacturer 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 the manufacturer, 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.

THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THESE 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.

BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

REFERENCE
86 A1 40EV 03