

# Hardware Information

Using the Virtual I/O Server

ESCALA POWER5



REFERENCE  
86 A1 24EW 00



# ESCALA POWER5

## Hardware Information

Using the Virtual I/O Server

### **Hardware**

July 2006

BULL CEDOC

357 AVENUE PATTON

B.P.20845

49008 ANGERS CEDEX 01

FRANCE

REFERENCE

86 A1 24EW 00

The following copyright notice protects this book under Copyright laws which prohibit such actions as, but not limited to, copying, distributing, modifying, and making derivative works.

Copyright © Bull SAS 1992, 2006

Printed in France

Suggestions and criticisms concerning the form, content, and presentation of this book are invited. A form is provided at the end of this book for this purpose.

To order additional copies of this book or other Bull Technical Publications, you are invited to use the Ordering Form also provided at the end of this book.

### **Trademarks and Acknowledgements**

We acknowledge the right of proprietors of trademarks mentioned in this book.

AIX® is a registered trademark of International Business Machines Corporation, and is being used under licence.

UNIX® is a registered trademark in the United States of America and other countries licensed exclusively through the Open Group.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries

# Table of Contents

<b>Using the Virtual I/O Server.....</b>	<b>1</b>
Printable PDFs.....	1
Virtual I/O Server overview.....	2
Virtual SCSI.....	2
Concepts for the Virtual I/O Server.....	3
Concepts for Virtual SCSI.....	3
Concepts for virtual networking.....	11
Virtual I/O Server management.....	13
Planning for the Virtual I/O Server.....	17
Specifications.....	17
Limitations and restrictions.....	18
Capacity planning.....	19
Redundancy considerations.....	29
Security considerations.....	33
Installing the Virtual I/O Server.....	33
Installing the Virtual I/O Server in an HMC environment.....	34
Installing the Virtual I/O Server in a non-HMC environment.....	39
Connecting to the Virtual I/O Server using OpenSSH.....	39
Configuration scenarios for the Virtual I/O Server.....	41
Scenario: Configuring a Virtual I/O Server without VLAN tagging.....	41
Scenario: Configuring a Virtual I/O Server using VLAN tagging.....	42
Scenario: Configuring shared Ethernet adapter failover.....	44
Scenario: Configuring Network Interface Backup in Virtual I/O clients without VLAN tagging.....	45
Scenario: Configuring Multi-Path I/O for AIX client logical partitions.....	47
Securing the Virtual I/O Server.....	49
Managing the Virtual I/O Server.....	52
Managing shared Ethernet adapters.....	52
Managing Virtual SCSI.....	63
Maintaining the Virtual I/O Server.....	71
Troubleshooting the Virtual I/O Server.....	77
Troubleshooting the Virtual I/O Server logical partition.....	77
Troubleshooting the client logical partition.....	81
Related information for the Virtual I/O Server.....	84
Virtual I/O Server command descriptions.....	84
Installation commands.....	85



## Using the Virtual I/O Server

The purpose of this information is to familiarize you with the Virtual I/O Server, to help you plan for the Virtual I/O Server in your computing environment, and to give you configuration and management instructions.

- **Printable PDFs**  
If you prefer a hardcopy version of this information, go here to print the PDF. Links to PDF documents about related topics are also included here.
  - **Virtual I/O Server overview**  
Learn the concepts of the Virtual I/O Server and its primary components.
  - **Concepts for the Virtual I/O Server**  
Become familiar with the Virtual I/O Server concepts, including the command-line interface, user types, virtual networking, and virtual SCSI.
  - **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.
  - **Installing the Virtual I/O Server**  
Find instructions for installing the Virtual I/O Server.
  - **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.
  - **Managing the Virtual I/O Server**  
Find information about managing Virtual I/O Server user types, adding and removing physical resources, and managing logical volumes. Also find information about backing up, restoring, updating, and monitoring the Virtual I/O Server.
  - **Troubleshooting the Virtual I/O Server**  
Find information about diagnosing Virtual I/O Server problems and information about how to correct those problems.
  - **Related information for the Virtual I/O Server**  
Find other information related to the Virtual I/O Server.
  - **Virtual I/O Server command descriptions**  
This topic includes descriptions of the Virtual I/O Server commands.
- 

## Printable PDFs

If you prefer a hardcopy version of this information, go here to print the PDF. Links to PDF documents about related topics are also included here.

To view or download the PDF version of this document, select [Using the Virtual I/O Server](#).

You can view or download these related topics:

- [Creating a virtual computing environment](#)
- [Virtual I/O Server Commands Reference](#)
- [Partitioning with Integrated Virtualization Manager](#)
- [Managing the Integrated Virtualization Manager](#)
- [Partitioning for AIX](#)
- [Partitioning for Linux](#)

Saving PDF files

To save a PDF on your workstation for viewing or printing:

1. Right-click the PDF in your browser (right-click the link above).
2. Click Save Target As... if you are using Internet Explorer. Click Save Link As... if you are using Netscape Communicator.
3. Navigate to the directory in which you would like to save the PDF.
4. Click Save.

## Downloading Adobe Acrobat Reader

You need Adobe Acrobat Reader to view or print these PDFs. You can download a copy from the [Adobe Web site](http://www.adobe.com/products/acrobat/readstep.html) (www.adobe.com/products/acrobat/readstep.html).

**Parent topic:** [Using the Virtual I/O Server](#)

---

## 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 resides in a POWER5 logical partition. This software facilitates the sharing of physical I/O resources between AIX and Linux 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 Advanced POWER Virtualization hardware feature.

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

- Sharing of physical resources between partitions on the system
- Creating partitions without requiring additional physical I/O resources
- Creating more partitions than there are I/O slots or physical devices available with the ability for partitions to have dedicated I/O, virtual I/O, or both
- Maximizing physical resource use 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:

- AIX 5.3 and later.
- SUSE LINUX Enterprise Server 9 for POWER
- Red Hat Enterprise Linux AS for POWER Version 3
- Red Hat Enterprise Linux AS for POWER Version 4

The Virtual I/O Server is comprised of 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 storage subsystem that provides standard SCSI-compliant logical unit numbers (LUNs). The Virtual I/O Server is capable of exporting a pool of heterogeneous physical storage as an 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 out 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:



- Disks backed by a logical volume
- Disks backed by a physical volume
- Optical devices (DVD-RAM and DVD-ROM)

For more information about virtual SCSI, see [Concepts for Virtual SCSI](#).

### Virtual networking

Shared Ethernet adapter allows logical partitions on the virtual local area network (VLAN) to share access to a physical Ethernet adapter and to communicate with systems and partitions outside the server. This function enables logical partitions on the internal VLAN to share the VLAN with standalone servers.

For more information about virtual networking, see [Concepts for virtual networking](#).

### Integrated Virtualization Manager

The Integrated Virtualization Manager provides a browser-based interface and a command-line interface that you can use to manage 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.

For more information about the Integrated Virtualization Manager, see [Managing the Integrated Virtualization Manager](#) and [Partitioning with the Integrated Virtualization Manager](#).

**Parent topic:** [Using the Virtual I/O Server](#)

---

## Concepts for the Virtual I/O Server

Become familiar with the Virtual I/O Server concepts, including the command-line interface, user types, virtual networking, and virtual SCSI.

- [Concepts for 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.
- [Concepts for virtual networking](#)  
Use this section to find information about virtual Ethernet, shared Ethernet adapter, shared Ethernet adapter failover, link aggregation, and VLAN.
- [Virtual I/O Server management](#)  
This topic contains information about Virtual I/O Server management interfaces, such as the Virtual I/O Server command-line interface and the Integrated Virtualization Manager. Virtual I/O Server user types are also explained

**Parent topic:** [Using the Virtual I/O Server](#)

---

## Concepts for 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 adapter 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
- Optical CD-ROM, DVD-RAM, and DVD-ROM

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 logical volumes. The logical volumes can then be assigned to different partitions. Therefore, virtual SCSI enables the sharing of adapters as well as disk devices. To make a physical or a logical volume available to a client 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 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.
- **Concepts for physical storage**  
Learn more about physical storage, logical volumes, and the devices and configurations that are supported by the Virtual I/O Server.
- **Concepts for 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.
- **Concepts for mapping devices**  
Mapping devices are used to facilitate the mapping of physical resources to a virtual device.

**Parent topic:** [Concepts for the Virtual I/O Server](#)

## 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. Although the SCSI LUNs are SCSI-compliant, they might not meet the needs of all applications, particularly those that exist in a distributed environment.

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, and logical volumes.

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 (vtscsiX). 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 recreating the physical-to-virtual allocations in a persistent manner when the Virtual I/O Server is restarted.

**Parent topic:** [Concepts for Virtual SCSI](#)

---

## Concepts for 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 might 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.
- **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.
- **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.

**Parent topic:** [Concepts for Virtual SCSI](#)

---

## Physical volumes

Physical volumes might 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](#).

The following commands are used to manage physical volumes:

Table 1. Physical volume commands and their descriptions

Physical volume command	Description
<a href="#">lspv</a>	Displays information about a physical volume within a volume group.
<a href="#">migratepv</a>	Moves allocated physical partitions from one physical volume to one or more other physical volumes.

Parent topic: [Concepts for physical storage](#)

---

## 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 discontinuous 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. For more information on volume groups, see [Volume groups](#).

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

Table 1. Logical volume commands and their descriptions

Logical volume command	Description
<a href="#">chlv</a>	Changes the characteristics of a logical volume.
<a href="#">cplv</a>	Copies the contents of a logical volume to a new logical volume.
<a href="#">extendlv</a>	Increases the size of a logical volume.
<a href="#">lslv</a>	Displays information about the logical volume.
<a href="#">mklv</a>	Creates a logical volume.
<a href="#">mklvcopy</a>	Creates a mirror of a logical volume.
<a href="#">rmlv</a>	Removes logical volumes from a volume group.
<a href="#">rmlvcopy</a>	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.

**Note:** Logical volumes used as virtual disks must be less than 1 TB in size. In addition, logical volumes on the Virtual I/O Server that are going to be used as virtual disks cannot be mirrored, striped, or have bad block relocation enabled.

- **Volume groups**  
Find information about volume groups.
- **Physical partitions**  
This topic contains information about physical partitions.
- **Logical partitions**  
This topic contains information logical storage partitions.
- **Quorums**  
Find information about quorums.
- **Storage pools**  
This topic includes overview information about storage pools.

**Parent topic:** [Concepts for physical storage](#)

---

## Volume groups

Find information about volume groups.

A volume group is a collection of 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.

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.

Some frequently used volume group commands are as follows:

Command	Description
<a href="#">mkvg</a>	Creates a volume group
<a href="#">extendvg</a>	Adds a physical volume to a volume group
<a href="#">chvg</a>	Changes the attributes of a volume group
<a href="#">reducevg</a>	Removes a physical volume from a volume group
<a href="#">lsvg</a>	Displays information about a volume group
<a href="#">exportvg</a>	Exports the definition of a volume group
<a href="#">importvg</a>	Imports a new volume group definition
<a href="#">syncvg</a>	Synchronizes logical volume copies that are not current
<a href="#">activatevg</a>	Activates a volume group
<a href="#">deactivatevg</a>	Deactivates a volume group

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. For more information, see the [migratepv](#) command description.

Parent topic: [Logical volumes](#)

---

## 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.

Parent topic: [Logical volumes](#)

---

## 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.

Parent topic: [Logical volumes](#)

---

## 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 LVM. 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.

Parent topic: [Logical volumes](#)

---

## Storage pools

This topic includes overview information about storage pools.

In Virtual I/O Server Version 1.2 and later, storage pools are available. Similar to volume groups, storage pools are collections of one or more physical volumes. The physical volumes that comprise a storage pool can be of varying sizes and types. Using storage pools, you are no longer required to have extensive knowledge on how to manage volume groups and logical volumes to create and assign logical storage to a client partition. Devices created using the storage pool are not limited to the size of the individual physical volumes.

Storage pools are created and managed using the following commands:

Table 1. Storage pool commands and their descriptions

Command	Description
<a href="#">mksp</a>	Creates a storage pool
<a href="#">chsp</a>	Changes the characteristics of a storage pool
<a href="#">lssp</a>	Displays information about a storage pool
<a href="#">mkdbsp</a>	Carves storage out of a pool and assigns it to a virtual SCSI adapter as a backing device
<a href="#">rmdbsp</a>	Disassociates a backing device from its virtual SCSI adapter and removes it from the system

There is a single default storage pool for each Virtual I/O Server partition that can be modified only by the prime administrator. Unless explicitly specified otherwise, the storage pool commands will operate on the default storage pool. This 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.

**Parent topic:** [Logical volumes](#)

---

## 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 might be backed by DVD drives. 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

The virtual optical device 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 partition and reassigned to the partition that will use the device. **Parent topic:** [Concepts for physical storage](#)

---

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

Parent topic: [Concepts for Virtual SCSI](#)

---

## 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. Two different 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

Regardless of whether the virtual SCSI disk is backed by a physical disk or a logical volume, 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 if a Virtual I/O Server is not responding to I/O requests. It is recommended that you enable this feature only in configurations in which devices are available to a client 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 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 partition error log. If one of the failed devices is the root volume group (rootvg) for the partition, and the rootvg is not available via another path or is not being mirrored on another Virtual I/O Server, the client 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.

**Parent topic:** [Concepts for virtual storage](#)

---

## 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 partitions. These are referred to as *virtual SCSI optical devices*. Virtual optical devices may be backed by DVD drives. 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

The virtual optical device 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 partition and reassigned to the partition that will use the device.

Virtual 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, or USB device.

**Parent topic:** [Concepts for virtual storage](#)

---

## Concepts for mapping devices

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

**Parent topic:** [Concepts for Virtual SCSI](#)

---

## Concepts for virtual networking

Use this section to find information about virtual Ethernet, shared Ethernet adapter, shared Ethernet adapter failover, link aggregation, and VLAN.

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

- **[Introduction to virtual networking](#)**  
This topic introduces virtual networking concepts. Each concept introduced will be discussed in more detail in subsequent sections.
- **[Virtual local area networks \(VLAN\)](#)**  
Virtual local area networks (VLAN) allows the physical network to be logically segmented.
- **[Virtual Ethernet adapters](#)**  
Virtual Ethernet adapters allow client logical partitions to send and receive network traffic without having a physical Ethernet adapter.
- **[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.

**Parent topic:** [Concepts for the Virtual I/O Server](#)

---

## Introduction to virtual networking

This topic introduces virtual networking concepts. Each concept introduced will be discussed in more detail in subsequent sections.

Virtual Ethernet technology allows logical partitions within the same system to communicate without having to use physical Ethernet adapters. Virtual Ethernet adapters are created by using the HMC and configured by using the Virtual I/O Server command-line interface. 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 or logical partitions. The Integrated Virtualization Manager can also be used to create and manage virtual Ethernet adapters. For more information, see [Configuring virtual Ethernet bridges on the managed system in the Partitioning with the Integrated Virtualization Manager](#) topic.

Logical partitions that require access to the external network must have either a dedicated Ethernet adapter or a virtual Ethernet adapter that sends and receives packets through a Shared Ethernet adapter on the Virtual I/O Server partition. Multiple client logical partitions can share this resource. The shared Ethernet adapter acts like a layer 2 bridge to forward outbound packets received from a virtual Ethernet adapter to the external network and forward inbound packets to the appropriate client logical partition over the virtual Ethernet link to that partition.

Consider using shared Ethernet on the Virtual I/O Server when the capacity or the bandwidth requirements of the individual partitions is inconsistent or is less than the total bandwidth of a physical Ethernet adapter. Partitions that use the full bandwidth or capacity of a physical Ethernet adapter should use dedicated Ethernet adapters.

The shared Ethernet adapter is configured by using the command-line interface on the Virtual I/O Server. Examples are shown later in this section.

The shared Ethernet adapter can be configured as a link aggregation device and with virtual local area network (VLAN) tagging. Link aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter. For example, multiple gigabit Ethernet adapters could be configured as a single IP address on the Virtual I/O Server. Link aggregation can also 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.

**Parent topic:** [Concepts for virtual networking](#)

---

## Virtual local area networks (VLAN)

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.

Parent topic: [Concepts for virtual networking](#)

---

## 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 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 that enable them to share a common logical network. The virtual Ethernet adapters are created and the VID assignments are done using the Hardware Management Console (HMC). The system transmits packets by copying the packet directly from the memory of the sender partition to the receive buffers of the receiver partition without any intermediate buffering of the packet.

Parent topic: [Concepts for virtual networking](#)

---

## 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.

While virtual Ethernet technology can provide communication between logical partitions on the same system, network access outside the system requires physical adapters. The Virtual I/O Server can provide network services to logical partitions without requiring each partition to own a physical device. The network access component of the Virtual I/O Server is called the *shared Ethernet adapter*.

A shared Ethernet adapter is a bridge between a physical Ethernet adapter or link aggregation and one or more virtual Ethernet adapters on the Virtual I/O Server. A shared Ethernet adapter enables logical partitions on the virtual Ethernet to share access to the physical Ethernet and communicate with stand-alone servers and logical partitions on other systems. The shared Ethernet adapter provides this access by connecting the internal VLANs with the VLANs on the external switches. This enables logical partitions to share the IP subnet with standalone systems and other external logical partitions. 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.

Parent topic: [Concepts for virtual networking](#)

---

## Virtual I/O Server management

This topic contains information about Virtual I/O Server management interfaces, such as the Virtual I/O Server command-line interface and the Integrated Virtualization Manager. Virtual I/O Server user types are also explained

- [Integrated Virtualization Manager](#)  
Use this topic to find information about the Integrated Virtualization Manager.
- [Virtual I/O Server command-line interface](#)  
Use this topic to find information about the Virtual I/O Server command-line interface.
- [User types](#)  
Use this topic to provide information about Virtual I/O Server user types and their user permissions.

Parent topic: [Concepts for the Virtual I/O Server](#)

---

## Integrated Virtualization Manager

Use this topic to find information about the Integrated Virtualization Manager.

The Integrated Virtualization Manager provides a graphical user interface available for use in environments where no HMC is present. You can use Integrated Virtualization Manager to perform system management tasks, such as creating partitions and assigning resources to those partitions. Where no HMC is present, the Integrated Virtualization Manager is installed and activated when the Virtual I/O Server is installed. Many of the Virtual I/O Server tasks can be performed using the Integrated Virtualization Manager.

**Note:** The Virtual I/O Server command-line interface is still accessible to users in a Integrated Virtualization Manager operating environment.

For more information about the Integrated Virtualization Manager, see [Partitioning with Integrated Virtualization Manager](#) and [Managing Integrated Virtualization Manager](#).

**Parent topic:** [Virtual I/O Server management](#)

---

## Virtual I/O Server command-line interface

Use this topic to find information about 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, LVM)
- Network configuration
- Software installation and update
- Security
- User management
- Maintenance tasks

In addition, in Integrated Virtualization Manager manager environments, the Virtual I/O Server command-line interface is used for partition management.

For a detailed description of each Virtual I/O Server command, see [Virtual I/O Server command descriptions](#).

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.

**Note:** For command reference information for the Virtual I/O Server, see [Virtual I/O Server command descriptions](#).

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

Table 1. Standard shell commands and their functions

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

For command reference information for the Virtual I/O Server, see [Virtual I/O Server command descriptions](#).

**Parent topic:** [Virtual I/O Server management](#)

---

## User types

Use this topic to provide information 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 is created by executing the [mkuser](#) command from the padmin user ID. There is no limit to the number of system administrator IDs that can be created. The system administrator user ID has access to all commands except the following:

- [cleargcl](#)
- [lsfailedlogin](#)
- [lsgcl](#)
- [mirrorios](#)
- [mkuser](#)
- [oem\\_setup\\_env](#)
- [rmuser](#)
- [shutdown](#)
- [unmirrorios](#)

### Service representative

The SR user is created to allow an service representative (SR) to log in to the system to perform diagnostic routines. SR users are created by running the [mkuser](#) command with the -sr flag.

Upon logging in, the SR user is placed directly into the diagnostic menus.

### Development engineer

The DE user is created to allow a development engineer (DE) to log in to the system and debug problems. DE users are created by running the **mkuser** command with the **-de** flag.

Parent topic: [Virtual I/O Server management](#)

---

## 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.
- **Limitations and restrictions**  
Find Virtual I/O Server configuration limitations.
- **Capacity planning**  
This topic includes capacity-planning considerations for the Virtual I/O Server, including information about hardware resources and limitations.
- **Redundancy considerations**  
Redundancy options are available at several levels in the virtual I/O environment. Multipathing and RAID redundancy options exist for both the Virtual I/O Server and client partitions. Ethernet link aggregation is also an option for the client 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.
- **Security considerations**  
Review the security considerations for Virtual SCSI, virtual Ethernet, and Shared Ethernet adapter and the additional security options available.

Parent topic: [Using 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 Advanced POWER Virtualization hardware feature is required. A logical partition with enough resources to share with other partitions is required. The following is a list of minimum hardware requirements that must be available to create the Virtual I/O Server:

Table 1. Resources that are required

Resource	Requirement
Hardware Management Console or Integrated Virtualization Manager	Either the HMC or Integrated Virtualization Manager is required to create the partition and assign resources.
Storage adapter	The server 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	At least 512 MB of memory is required.
Processor	At least .1 processor is required.

The following table defines the limitations for storage management.









































































































































































