System i and System p

Using the Virtual I/O Server

System i and System p

Using the Virtual I/O Server

Note

Before using this information and the product it supports, read the information in "Notices" on page 179 and the *IBM Systems Safety Information* manual, G229-9054.

Ninth Edition (September 2007)

This edition applies to IBM Virtual I/O Server version 1.4 and to all subsequent releases until otherwise indicated in new editions.

Contents

About this topic

This topic provides system operators with information about installing, configuring, managing, and using the Virtual I/O Server.

For information about the accessibility features of this product, for users who have a physical disability, see “Accessibility features,” on page 177.

Using the Virtual I/O Server

The Virtual I/O Server is software that facilitates the sharing of physical I/O resources between client logical partitions within the server. Learn about the Virtual I/O Server and how to use it in your computing environment.

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.

The following information applies to Virtual I/O Server version 1.3 and earlier. For information pertaining to Virtual I/O Server version 1.4 and later, see the PowerVM™ Editions Operations Guide .

What's new for Using the Virtual I/O Server

Learn about new and significantly changed information for the Virtual I/O Server version 1.4.

GARP VLAN Registration Protocol (GVRP) on Shared Ethernet Adapters

You can enable or disable GVRP on your Shared Ethernet Adapters to control dynamic registration of virtual LANs over networks.

- "Enabling and disabling GVRP" on page 100
- "GARP VLAN Registration Protocol statistics" on page 163
- "Shared Ethernet Adapter failover" on page 48
- "Shared Ethernet Adapters" on page 13
- "Viewing statistics for Ethernet drivers and devices" on page 147

LDAP client

You can configure Virtual I/O Server version 1.4 as an LDAP client. Then you can manage Virtual I/O Server from an LDAP server.

- "Configuring the Virtual I/O Server as an LDAP client" on page 126
- "Managing users on the Virtual I/O Server" on page 118
- ldapadd Command
- ldapsearch Command
- mkldap Command

New HMC interface

HMC version 7 displays a different interface than prior versions of the HMC. The following tasks are new or updated to include how to use the new interface to complete the procedures that require an HMC:

- "Entering the activation code for Advanced POWER Virtualization using the HMC version 7" on page 51
- "Creating a Shared Ethernet Adapter using HMC version 7" on page 95
- "Creating the Virtual I/O Server logical partition and partition profile using HMC version 7" on page 61
- "Installing the Virtual I/O Server from CD or DVD" on page 63
- "Installing the Virtual I/O Server manually using the HMC version 7" on page 60

- “Recovering from disks not displaying in SMS” on page 172

SNMP

You can enable, disable, and work with SNMP on the Virtual I/O Server.

- “Managing SNMP on the Virtual I/O Server” on page 101
- cl_snmp Command
- snmp_info Command
- snmp_trap Command

System plans

You can use the System Planning Tool (SPT), version 2.0, and the Deploy System Plan wizard on the Hardware Management Console (HMC), version 7, to plan for and deploy the Virtual I/O Server. In addition to basic configuration, like creating virtual Ethernet and virtual SCSI adapters, you can now plan for and deploy advanced configuration options like Shared Ethernet Adapter failover and multi-path I/O.

- “Installing the Virtual I/O Server by deploying a system plan” on page 51
- “Planning for Virtual I/O Server using system plans” on page 29
- “System plan overview for the HMC” on page 20

Tivoli® agents and clients

Virtual I/O Server version 1.4 includes the IBM® Tivoli Monitoring premium agent, the IBM Tivoli Usage and Accounting Manager agent, and the IBM Tivoli Storage Manager client. The following information is new or changed to help you work with the agents and clients:

- “Backing up the Virtual I/O Server using IBM Tivoli Storage Manager” on page 136
- “Configuration attributes for IBM Tivoli agents and clients” on page 124
- “Configuring the IBM Tivoli Monitoring agent” on page 119
- “Configuring the IBM Tivoli Storage Manager client” on page 123
- “Configuring the IBM Tivoli Usage and Accounting Manager agent” on page 121
- “IBM Tivoli agents and clients on the Virtual I/O Server” on page 28
- “Restoring the Virtual I/O Server using IBM Tivoli Storage Manager” on page 144
- cfgsvc Command
- dsmc Command
- lssvc Command
- startsvc Command
- stopsvc Command

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 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 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 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:

- AIX 5.3 or later
- SUSE Linux Enterprise Server 9 for POWER (or later)
- Red Hat Enterprise Linux AS for POWER Version 3 (update 2 or later)
- Red Hat Enterprise Linux AS for POWER Version 4 (or later)

For the most recent information about devices that are supported on the Virtual I/O Server, to download Virtual I/O Server fixes and updates, and to find additional information about the Virtual I/O Server, see the Virtual I/O Server Web site.

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 an 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:

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

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 stand-alone servers.

Integrated Virtualization Manager

The Integrated Virtualization Manager provides a browser-based interface and a command-line interface that you can use to manage IBM System p5® and IBM eServer™ pSeries® servers that use the IBM 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.

Related concepts

"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" on page 12
Use this section to find information about virtual Ethernet, Shared Ethernet Adapter, Shared Ethernet Adapter failover, Link Aggregation (or EtherChannel), and VLAN.

Related tasks

Managing the Integrated Virtualization Manager

Partitioning with the Integrated Virtualization Manager

Related information

Virtual I/O Server Support for UNIX servers and Midrange servers

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.

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.

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 (vtscsi*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.

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 it's 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
lspv	Displays information about a physical volume within a volume group.
migratepv	Moves allocated physical partitions from one physical volume to one ore more other physical volumes.

Related tasks

"Identifying exportable disks" on page 111
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).

Related reference

lspv Command

migratepv Command

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 2. 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.

Table 2. Logical volume commands and their descriptions (continued)

Logical volume command	Description
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 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.

Related reference

chlv Command

cplv Command

extendlv Command

lslv Command

mklv Command

mklvcopy Command

rmlv Command

rmlvcopy Command

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.

Table 3. 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 3. 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.

Related concepts

"Physical partitions"
This topic contains information about physical partitions.

Related reference

activatevg Command

chvg Command

deactivatevg Command

exportvg Command

extendvg Command

importvg Command

lsvg Command

migratepv Command

mkvg Command

reducevg Command

syncvg Command

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.

Related reference

activatevg Command

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 mange 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 4. Storage pool commands and their descriptions

Command	Description
chsp	Changes the characteristics of a storage pool
lssp	Displays information about a storage pool
mkdbsp	Carves storage our of a pool and assigns it to a virtual SCSI adapter as a backing device
mksp	Creates a storage pool
rmdbsp	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.

Related reference

chsp Command

lssp Command

mkdbsp Command

mksp Command

rmdbsp Command

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.

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.

The Virtual I/O Server might virtualize, or export, disks and optical devices, such as CD-ROM drives and DVD drives, as virtual devices. For a list of supported disks and optical devices, see the datasheet available on the Virtual I/O Server Support for UNIX® servers and Midrange servers Web site. For information about configuring virtual SCSI devices, see Creating the virtual target device on the Virtual I/O Server.

Related tasks

"Creating the virtual target device on the Virtual I/O Server" on page 112
Find instructions for creating a virtual target device on the Virtual I/O Server.

Related information

Virtual I/O Server Support for UNIX servers and Midrange servers

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. 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.

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.

Concepts for mapping devices

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

Concepts for virtual networking

Use this section to find information about virtual Ethernet, Shared Ethernet Adapter, Shared Ethernet Adapter failover, Link Aggregation (or EtherChannel), 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.

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 that enable them to share a common logical network. 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.

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.

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 requirements of the individual partitions is inconsistent with, 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.
- When you need an Ethernet connection, but there is no slot available in which to install a dedicated adapter.

Related tasks

Configuring virtual Ethernet bridges on the managed system

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.

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.

The Shared Ethernet Adapter enables logical partitions on the virtual network to share access to the physical network and communicate with stand-alone servers and logical partitions on other systems. It 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. This enables logical partitions to share the IP subnet with standalone 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 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 (GVRP)

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

With Virtual I/O Server version 1.4, you can assign a Logical Host Ethernet port, of a Logical Host Ethernet Adapter (LHEA), 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 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.

Consider assigning a Shared Ethernet Adapter to an 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 adapters in the future, then you can use the ports of the LHEA for network connectivity rather than the Shared Ethernet Adapter.

Related concepts

"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).

"Shared Ethernet Adapter failover" on page 48
Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption.

"Planning for Shared Ethernet Adapters" on page 35
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.

Related tasks

"Configuring a Shared Ethernet Adapter" on page 97
Find instructions for configuring Shared Ethernet Adapters.

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 `ent`*`X`* on AIX logical partitions, `CMN`*`XX`* on i5/OS® logical partitions, and `eth`*`X`* 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. Because you are bridging the Ethernet connection through the Virtual I/O Server, the connection might not perform as well as a logical network. 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 version 7 or later of 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.

Related concepts

"Shared Ethernet Adapters" on page 13
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.

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, which enables them to act as a single Ethernet device. It 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 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

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 AIX and Linux logical partitions on a single IBM System p® server. On OpenPower® servers, the Integrated Virtualization Manager supports only Linux logical partitions.

The Integrated Virtualization Manager is supported only on specific server models.

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.

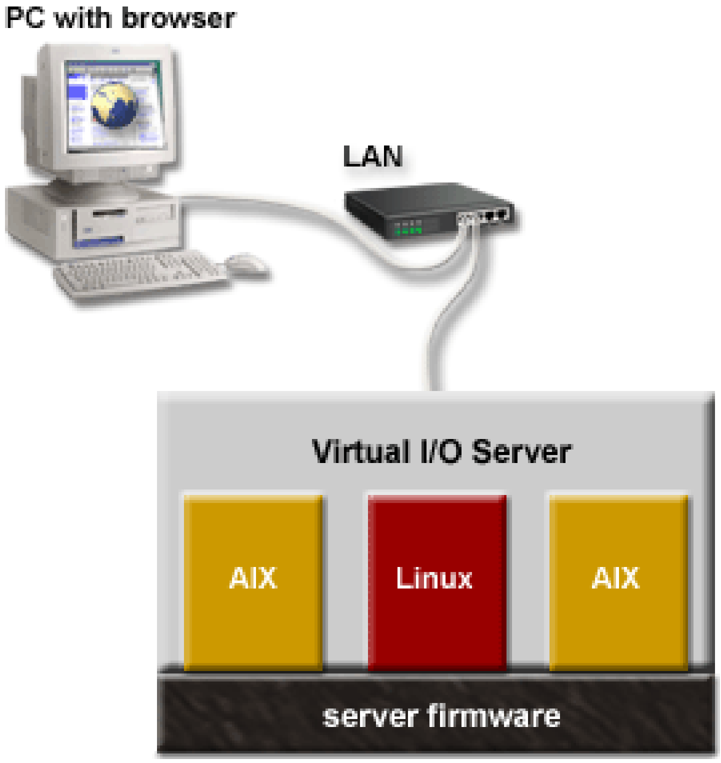

This figure illustrates Virtual I/O Server in its own logical partition, and the AIX and Linux logical partitions that are managed by the Virtual I/O Server logical partition. The browser on the PC connects to the Integrated Virtualization Manager interface over a network, and you can use the Integrated 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

Related tasks

Managing the Integrated Virtualization Manager

Partitioning with 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.

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.

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 5. 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.

Table 5. Standard shell commands and their functions (continued)

Command	Function
vi	Edits files with full screen display.
wc	Counts the number of lines, words, and bytes or 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 (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`

Related reference

awk Command

cat Command

chmod Command

cp command

date Command

errlog Command

grep Command

ls Command

lsgcl Command

man Command

mkdir Command

more Command

rm Command

sed Command

stty Command

tee Command

vi Command

wc Command

who Command

System plan overview for the HMC

Learn about system plan concepts and operations, as well as understand the high-level tasks that you can perform with system plans when using the Hardware Management Console (HMC).

A *system plan* is a specification of the hardware and the logical partitions contained in one or more systems. A system plan is stored in a *system-plan file*, which has a file suffix of .sysplan. A system-plan file can contain more than one system plan, although multiple plans in a single file are not common. After you create a system plan, you also can also view, delete, and export the system plan.

System plans have a number of valuable uses. For example, you can use system plans to accomplish the following goals:

- You can create a system plan as a means of capturing up-to-date system documentation. The system plan provides a record of the hardware and partition configuration of the managed system at a given time.
- You can use a system plan that you create for system documentation as part of your disaster recovery planning. You can export the system-plan file to an offsite location or to removable media for offsite storage so that you have the system documentation that you need available to you if you must recover a managed system.
- You can use system plans as audit records to track system resources for accounting and accountability purposes by exporting them to a spreadsheet.
- You can use system plans to help you plan new workloads that require additional system and hardware resources. You can use a system plan, along with appropriate capacity planning information, to make decisions about whether your current system can handle a new workload.
- You can deploy this system plan to other systems that this HMC manages that have hardware that is identical to the hardware in the system plan. In this way, you can rapidly configure and use other, similar systems in your business.
- You can export the system plan to another HMC and use it to deploy the system plan to other systems that the target HMC manages that have hardware that is identical to the hardware in the system plan. In this case and the previous case, you can use the system plan to create logical partitions on new managed systems that do not already have logical partitions created on them.

To create logical partitions from a system plan, you must first complete the following tasks:

1. Create the system plan.
2. Import the system plan (when necessary).
3. Deploy the system plan.

After you create a system plan, you also can also view, delete, and export the system plan. The following table provides a complete overview of system plan tasks.

Table 6. Overview of the tasks for system plans

Task	Overview
Create a system plan	You can create system plans by using any of the following methods: • System Planning Tool (SPT) *SPT* helps you design a system to fit your needs, whether you want to design a logically partitioned system or to design an unpartitioned system. SPT incorporates the function from Workload Estimator (WLE) to help you create an overall system plan. The SPT opens the WLE to help you gather and integrate workload data, and provides advanced users with the option of creating a system plan without the help of additional tools. To help you get started, SPT provides the following options: – You can use the sample system plans that SPT provides as a starting point for planning your system – You can create a system plan based on existing performance data – You can create a system plan based on new or anticipated workloads – You can export a system plan as a .cfr file and import it into the marketing configurator (eConfig) tool to use for ordering a system. When you import the .cfr file into the eConfig tool, the tool populates your order with the information from the .cfr file. However, the .cfr file does not contain all the information that the eConfig tool requires and you will need to enter all required information before you can submit your order. • Hardware Management Console (HMC) Web user interface You can use the HMC to create a system plan based on the configuration of one managed system and can use the HMC to deploy that plan to another managed system. Based on the logical partition configuration in the system plan, the HMC creates logical partitions on the managed system to which it deploys the system plan. Depending on the contents of the system plan, the HMC can install operating environments on the partitions in the plan and, if the plan contains Virtual I/O Server provisioning information for a partition, such as storage assignments, the HMC can make these resource assignments for the partition. • HMC command-line interface You also can use the **mksysplan** command to create a system plan. After the system plan is created, you can also use the command-line interface to deploy that plan to a managed system. Based on the logical partition configuration in the system plan, the HMC creates logical partitions on the managed system to which it deploys the system plan.
Import the system plan	Before you can use a system plan to create logical partitions, the system-plan file must exist on the HMC that manages the managed system to which you want to deploy the system plan. If the system-plan file does not already exist on the HMC, you must import the file into the HMC. You can use the HMC Web user interface to import the file into the HMC from one of the following sources: • Upload the system-plan file from the remote console (the computer from which you remotely access the HMC) • Copy the system-plan file to media (optical disc or USB drive), insert the media into the HMC, and import the file from the media. • Download the system-plan file from a remote FTP site. **Note:** You can also use the HMC command-line interface to import a system plan. After you import the system-plan file into an HMC, you can deploy the system plan within that file to other systems that the HMC manages.

Table 6. Overview of the tasks for system plans (continued)

Task	Overview
Deploy the system plan	You can choose to deploy a system plan in stages, with some logical partitions being created in one stage, and other logical partitions being created in later stages. You cannot, however, deploy a system plan to a managed system if the managed system already has logical partitions. The managed system must be in the manufacturing default configuration. Also, if you want to deploy a system plan in stages, you need to create a new system plan if you change the resource allocations on the logical partitions on the managed system between stages to avoid validation problems in later stages. When you deploy a system plan by using the HMC Web user interface, the HMC validates the system plan. The managed system on which you deploy a system plan must have hardware that is identical to the hardware in the system plan. The HMC deploys a system plan to a managed system only if the system plan level is supported by the HMC, the format of the system plan is valid, and the hardware and each existing logical partition on the managed system passes validation. If the system plan contains installation information about the Virtual I/O Server, you can use the Deploy System Plan wizard to install the Virtual I/O Server and assign virtual networking and storage resources for the client logical partitions.
Export the system plan	You can use the HMC Web user interface to export a system-plan file from the HMC to one of the following locations: • Save the system-plan file to the remote console (the computer from which you remotely access the HMC). • Export the system-plan file to media that is mounted to the HMC (such as optical discs or USB drives). • Download the system-plan file to a remote FTP site. **Note:** You can also use the HMC command-line interface to export a system plan.
View the system plan	You can look at the contents of a system-plan file in the HMC by using the System Plan Viewer that is integrated with the HMC. 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 displaying EADS boundary lines. You can open a system plan in the System Plan Viewer, either by using the View System Plan task or by clicking the name of a system plan. When you start the System Plan Viewer, you must enter your HMC user ID and password before you can view the system plan.
Print the system plan	You can use the System Plan Viewer to print a system plan that you have open in the Viewer. You can print all of the system plan or a portion of the system plan, depending on the current view of the system plan. To print the current view of the system plan, click **Print** in the Actions pane of the System Plan Viewer.
Delete the system plan	You can delete unnecessary system plans from your HMC.

Optimizing system plan hardware information

The amount of hardware information that the HMC can capture in a new system plan varies based on the method that the HMC uses to gather the hardware information. Setting up your environment to maximize inventory gathering allows the HMC to capture more complete information about the hardware allocated to the partitions on the managed system. For example, the HMC can capture disk drive and tape drive configuration information for an active partition in the new system plan. However, doing so can cause system plan creation to take several more minutes to complete.

There are two methods that the HMC potentially can use:

- Inventory gathering, which is available for HMC Version 7 Release 3.1.0 and later
- Hardware discovery, which is available for some systems with HMC Version 7 Release 3.2.0 and later

System plan inventory gathering

The HMC always performs inventory gathering to capture detailed information for hardware that has an assignment to an active partition.

Note: Beginning with HMC Version 7.3.2, you can use the hardware discovery process to gather information about hardware assignments for an inactive partition or hardware on a managed system that does not have a partition assignment.

To optimize the amount of and type of hardware information that the inventory-gathering process is able to capture, ensure that you meet the following prerequisites and conditions:

- You must set up Resource Monitoring and Control (RMC) prior to creating a system plan. Using RMC ensures that the inventory-gathering process can capture more detailed hardware information. Without RMC, the inventory-gathering process is not able to detect the types of disk drives installed on a managed system.

  Note: i5/OS partitions respond to RMC requests from the HMC by means of the Management Server.
- To ensure that Linux systems and partitions can perform inventory gathering, you must load the IBM Installation Toolkit for Linux on POWER, which is available at the IBM Service and productivity tools Web site (http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/installtools/home.html).
- You must have the managed system in the ’Standby’ state or you must power on the managed system and activate the logical partitions on the managed system before creating the system plan.

Note: It is possible for a partition to have more than one HMC set up to manage it. In this situation, if the partition is an i5/OS partition and you want to use RMC to create a new system plan, ensure that you create the system plan from the primary HMC for the partition because redundant HMCs cannot use RMC.

System plan hardware discovery

In some cases, the HMC Version 7.3.2 can use hardware discovery, in addition to the inventory-gathering process, to capture hardware information for a new system plan. Using hardware discovery, you can capture information about hardware that does not have a partition assignment, as well as hardware with assignments to inactive partitions.

On a system that can use hardware discovery, the hardware discovery process runs whenever the system is powered on in *hardware discovery* mode. The hardware discovery process writes hardware inventory information to a cache on the system. The hardware inventory cache ensures that a certain amount of hardware information is available on the system when you create a system plan. The HMC can capture the information in the cache for a system plan when partitions are active and the HMC cannot perform fresh hardware discovery on the partition.

Note: It is recommended that you power on the system in hardware discovery mode whenever you add or change hardware on the system.

If the managed system is capable of hardware discovery, the Create System Plan page provides an additional option that you can select to capture a broader range of hardware information for the new system plan. This option, called **Retrieve inactive and unallocated hardware resources**, allows you to capture hardware configuration information for the managed system, regardless of the state of the hardware.

When you create a system plan and do not select the **Retrieve inactive and unallocated hardware resources** option, the HMC does not perform hardware discovery. The HMC still performs inventory gathering and retrieves hardware information for any active partitions on the managed server. The resulting new system plan contains hardware information from the inventory-gathering process, as well as hardware information from the hardware inventory cache on the system.

To use the hardware discovery process, ensure that you meet the following prerequisites and conditions:

Available processing capability:
The hardware discovery process requires a minimum .5 processor be available for it to use.

Memory capability:
The hardware discovery process requires a minimum of 256 MB of free memory for it to use.

Partition state:
To maximize the information that the hardware discovery process can capture, partitions on the managed server must be inactive. If a partition is active, the hardware discovery process cannot capture fresh information from the partition and instead retrieves information about the hardware assigned to the inactive partition from the hardware inventory cache on the managed system.

By setting up your system to optimize the hardware information that you capture in a system plan that you create by using the HMC, you ensure that your system plan provides you with the most valuable information possible. It also ensures that you have the most usable configuration information possible when you convert the system plan for use in the System Planning Tool (SPT). The following table describes the type of hardware information that you can expect to see in a system plan that you convert, based on the system management product that you use to create the plan.

Table 7. Type of hardware information available in a system plan that you create in the HMC and convert to use in the SPT

Expected Conversion Results			
Partition	**HMC version 7 release 3.1.0 and earlier**	**HMC version 7 release 3.2.0 and later**	**Integrated Virtualization Manager**
i5/OS	Most cards. No disk, tape, CD, SCSI.	More cards. Some disk. No tape, CD, SCSI.	Not applicable.
All other operating systems	Very few cards. No disk, tape, CD, SCSI.	Most cards. Some disk. No tape, CD, SCSI.	Few, if any, cards. No disk, tape, CD, SCSI.

System plan validation

When validating the hardware on the managed system, the HMC compares the following information from the system plan with the hardware available on the managed system:

- Processor and memory amounts, including 5250 commercial processing workload (5250 CPW) where applicable
- 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.

For example, 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. A server that 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 adapters would allow the system to pass validation. A server that contains 4 GB of memory can cause the system to fail validation. A system plan can also fail 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 HMC does not validate the disk drives that are attached to physical I/O adapters against the disk drives specified in the system plan. You must ensure that the disk drives installed in the managed system support your desired logical partition configuration. Embedded devices automatically pass hardware validation because they are embedded into the system and cannot be removed.

If any step fails, validation fails for the existing logical partition. Any existing partition found on the managed system must appear in the system plan and must match the system plan as it appears in the managed system. For example, hardware on the managed system must at least match the hardware specified in the system plan. When validating an existing logical partition, the HMC validates the following items for that logical partition:

1. Is there a logical partition in the system plan that has the same partition ID as the existing logical partition specified in the machine default configuration?
2. Does the existing logical partition have partition profiles that match each partition profile specified for the logical partition in the system plan?
3. Do the partition profiles for the existing logical partitions contain the resources specified in the corresponding partition profiles in the system plan?

For example, if the server has an existing logical partition with a partition ID of 1, the HMC examines the logical partition in the system plan that has a partition ID of 1. If this logical partition exists and has a partition profile that is named SUPPORT, the HMC looks at the existing logical partition to see if it also has a partition profile that is named SUPPORT. If so, the HMC verifies that the resources specified in the SUPPORT partition profile in the system plan are contained in the SUPPORT partition profile in the existing logical partition.

When the HMC validates partition profiles, it compares the following resources in the partition profiles:

- Processor and memory amounts, including 5250 commercial processing workload (5250 CPW) where applicable
- Physical I/O slot assignments

The following examples illustrate how the HMC compares resources in the partition profiles during the validation process to determine whether the system plan is valid for a managed system:

- If the SUPPORT partition profile in the system plan specifies 2 GB of memory and the SUPPORT partition profile for the existing logical partition specifies 3 GB of memory, the amount of memory is valid.
- If the SUPPORT partition profile in the system plan specifies 4 GB of memory and the SUPPORT partition profile for the existing logical partition specifies 3 GB of memory, the amount of memory is not valid.
- If physical I/O slot P1 is assigned to the SUPPORT partition profile in the system plan but not to the SUPPORT partition profile for the existing logical partition, the physical slot assignment is not valid.
- If physical I/O slot P2 is not assigned to the SUPPORT partition profile in the system plan, it does not matter whether slot P2 is assigned to the SUPPORT partition profile for the existing logical partition.

If the system plan contains installation information for the Virtual I/O Server, you can use the Deploy System Plan wizard to install the Virtual I/O Server and to set up virtual networking and storage resources for the client logical partitions of the Virtual I/O Server.

Note: The HMC cannot install AIX or Linux or i5/OS operating environments on logical partitions.

Related tasks

"Planning for Virtual I/O Server using system plans" on page 29
You can use the System Planning Tool (SPT) to create a system plan that includes configuration specifications for the Virtual I/O Server. You can also use the Hardware Management Console (HMC) or the Integrated Virtualization Manager to create a system plan based on an existing system configuration.

"Installing the Virtual I/O Server by deploying a system plan" on page 51
When you deploy a system plan that includes the Virtual I/O Server, 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.

User types for the Virtual I/O Server

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 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 IBM 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 IBM development engineer can log in to the system and debug problems.

Related tasks

"Managing users on the Virtual I/O Server" on page 118
Find commands for creating, listing, changing, switching, and removing users.

Related reference

lsfailedlogin Command

lsgcl Command

mirrorios Command

mkuser Command

oem_setup_env Command

rmuser Command

shutdown Command

unmirrorios Command

IBM Tivoli agents and clients on the Virtual I/O Server

Learn about the IBM Tivoli Monitoring agent, IBM Tivoli Storage Manager client, and the IBM Tivoli Usage and Accounting Manager agent packaged with the Virtual I/O Server.

IBM Tivoli Monitoring

Virtual I/O Server V1.3.0.1 (fix pack 8.1), includes the IBM Tivoli Monitoring System Edition for System p agent. IBM Tivoli Monitoring System Edition for System p enables you to monitor the health and availability of multiple IBM System p servers (including the Virtual I/O Server) from the Tivoli Enterprise Portal. IBM Tivoli Monitoring System Edition for System p 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 IBM Tivoli Monitoring.

IBM Tivoli Storage Manager

Virtual I/O Server 1.4 includes the IBM 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, including IBM System p servers. Configuring the Tivoli Storage Manager client on the Virtual I/O Server enables you to include the Virtual I/O Server in your standard backup framework.

IBM Tivoli Usage and Accounting Manager

Virtual I/O Server 1.4 includes the IBM Tivoli Usage and Accounting Manager agent on the Virtual I/O Server. IBM 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. IBM 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, i5/OS, and VMware.

Related tasks

"Configuring the IBM Tivoli agents and clients on the Virtual I/O Server" on page 119
You can configure and start the IBM Tivoli Monitoring agent, IBM Tivoli Usage and Accounting Manager agent, and the IBM Tivoli Storage Manager client.

Related information

ITM 6.1 documentation

IBM Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide

IBM Tivoli Usage and Accounting Manager Data Collectors for UNIX and Linux User's Guide

Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide

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.

Planning for Virtual I/O Server using system plans

You can use the System Planning Tool (SPT) to create a system plan that includes configuration specifications for the Virtual I/O Server. You can also use the Hardware Management Console (HMC) or the Integrated Virtualization Manager to create a system plan based on an existing system configuration.

SPT is a PC-based browser application that can assist you in planning and designing a new system. SPT validates your plan against system requirements and prevents you from exceeding system recommendations. It also incorporates the IBM Systems Workload Estimator (WLE) to help you plan for workloads and performance. The output is a system-plan file that you can deploy to a managed system.

With SPT version 2.0 and later, you can include the following configuration specifications for the Virtual I/O Server in your system plan:

- Backup virtual Ethernet adapters
- EtherChannel or Link Aggregation devices
- Mirroring
- Multi-path I/O
- SAN volumes
- Shared Ethernet Adapter failover
- Shared Ethernet Adapters
- Storage pools
- Virtual Ethernet adapter mappings between the Virtual I/O Server and its client logical partitions
- Virtual Ethernet adapters
- Virtual LANs
- Virtual SCSI adapter mappings between the Virtual I/O Server and its client logical partitions
- Virtual SCSI adapters

SPT currently does not help you plan for high availability on client logical partitions or Redundant Array of Independent Disks (RAID) solutions for the Virtual I/O Server.

To create a system plan that includes configuration specifications for the Virtual I/O Server, complete one of the following tasks:

- Create a system plan using SPT. For instructions, see the System Planning Tool Web site.

- Create a system plan based on an existing system configuration. For instructions, see Creating a system plan from an existing system configuration using HMC version 7 or Creating a system plan from an existing system configuration using the Integrated Virtualization Manager.

  Alternatively, you can use the mksysplan command to create a system plan based on an existing system configuration. The mksysplan command is available from the HMC and the Integrated Virtualization Manager.

After you have created a system plan, you can deploy the system plan to the managed system. System plans can be deployed to a system managed by the HMC or the Integrated Virtualization Manager. For instructions, see Deploying a system plan using HMC version 7 or Deploying a system plan using the Integrated Virtualization Manager.

On systems managed by the HMC, the HMC must be at version 7 or later. When you deploy the system plan, the HMC automatically creates the Virtual I/O Server logical partition and partition profile and installs the Virtual I/O Server based on the configuration specifications in the system plan.

On systems managed by the Integrated Virtualization Manager, the Integrated Virtualization Manager must be at version 1.4 or later. When you deploy the system plan, the Deploy System Plan wizard configures the management partition based on the configuration specification in the system plan.

System plans 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, the following items cannot be configured on the managed system before you deploy a system plan:

- Client logical partitions
- Virtual SCSI adapters
- Virtual Ethernet adapters
- Shared Ethernet Adapters
- EtherChannel adapters, or Link Aggregation devices
- Storage pools
- Backing devices

If you try to deploy a system plan on a system with any of these devices already configured, the Deploy System Plan wizard will fail the validation step.

Because of this limitation, the Deploy System Plan wizard will not allow you to partially deploy the Virtual I/O Server logical partition.

Related concepts

"System plan overview for the HMC" on page 20
Learn about system plan concepts and operations, as well as understand the high-level tasks that you can perform with system plans when using the Hardware Management Console (HMC).

"High Availability Cluster Multi-Processing" on page 46
Learn about High Availability Cluster Multi-Processing (HACMP™) in the Virtual I/O Server.

"RAID" on page 49
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.

Related tasks

Deploying a system plan using the Integrated Virtualization Manager

"Deploying a system plan by using HMC Version 7" on page 54
You can use the Hardware Management Console (HMC) to deploy all or part of a system plan to a managed system.

Related reference

HMC mksysplan Command

IVM mksysplan Command

Related information

System Planning Tool

Creating a system plan by using HMC Version 7

You can use the Hardware Management Console (HMC) Version 7 to create a new system plan based on an existing system configuration, and then deploy that system plan to other managed systems.

You can use the newly-created system plan to create identical logical partition configurations on managed systems with identical hardware. The new system plan contains specifications for the logical partitions and partition profiles of the managed system that you used to create the plan. The new system plan also can contain hardware information that the HMC is able to obtain from the selected managed system. However, the amount of hardware information that the HMC can capture for the new system plan varies based on the method that the HMC uses to gather the hardware information. There are two methods that the HMC potentially can use: inventory gathering and hardware discovery. For example, using inventory gathering, the HMC can detect virtual device configuration information for the Virtual I/O Server. Additionally, the HMC can use one or both of these methods to detect disk and tape information for i5/OS.

Inventory gathering prerequisites

The HMC always performs inventory gathering to capture detailed information for hardware that has an assignment to an active partition. To maximize the amount of data that the inventory gathering process of the HMC is able to collect from the managed system, ensure that you complete the following tasks:

- Ensure that the managed system in the 'Standby' state or that the managed system is powered on.

  Note: You cannot create a system plan if the managed server is in either the power off state or the recovery state.
- Ensure that all the logical partitions on the managed system from which you plan to base the new system plan are activated.
- To ensure that Linux systems and partitions can perform inventory gathering, you must load the IBM Installation Toolkit for Linux on POWER, which is available at the IBM Service and productivity tools Web site (http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/installtools/home.html).

- Ensure that there is an Resource Monitoring and Control (RMC) connection between the HMC and each logical partition. An RMC connection is required for the inventory-gathering process. An RMC connection also is required to configure Virtual I/O Server and to collect data for Virtual I/O Server device mappings.

  Note: It is possible for an i5/OS partition to have more than one HMC set up to manage it. In this situation, if you want to use RMC to create a new system plan, you must ensure that you create the system plan from the primary HMC for the partition because secondary HMCs cannot use RMC.

To ensure that the HMC can use RMC, complete the following steps:

1. In the HMC navigation pane, select **HMC Management**.
2. In the contents pane, select **Change Network Settings** to display the Customize Network Settings window.
3. Click **LAN Adapters**, select the appropriate adapter from the list, and click **Details**.
4. On the LAN Adapter page of the LAN Adapters Details window, ensure that **Partition communication** is selected.
5. On the Firewall page, in the Firewall Settings list, select all instances of RMC, and click **Allow Incoming**, if necessary.
6. Click **OK** to close the LAN Adapter Details window.
7. Click **OK** to close the Customize Network Settings window.
8. Restart the HMC if you made any changes to these configuration settings.

For some operating systems, you might need to perform additional steps to ensure that RMC is configured and running correctly. To learn more about configuring and using RMC, review the appropriate operating system documentation.

Hardware discovery prerequisites

If the managed system supports hardware discovery, the HMC can use it, in addition to the inventory-gathering process, to capture hardware information for a new system plan. The hardware discovery process allows you to capture hardware configuration information, regardless of the state of the hardware. Using hardware discovery, you can capture information about hardware that does not have a partition assignment, as well as hardware with assignments to inactive partitions.

To use the hardware discovery process, ensure that you complete the following tasks:

- Ensure that there is a minimum of .5 processor available.
- Ensure that there is a minimum of 256 MB of free memory available.
- Ensure that all partitions on the managed server for which you want to use the hardware discovery process are inactive. If a partition is active, the hardware discovery process cannot capture fresh information from the partition and retrieves information about the hardware assigned to the inactive partition from the hardware inventory cache on the managed system instead.

  Note: Hardware discovery does not require the use of RMC.

To create a system plan from Version 7 of the Hardware Management Console, complete the following steps:

1. From the navigation area, select **System Plans**. The System Plans page opens.
2. From the Tasks area, select **Create System Plan**. The Create System Plan window opens.
3. Select the managed system that you want to use as the basis for the new system plan.
4. Enter a name and description for the new system plan.

5. Optional: Select whether you want to retrieve inactive and unallocated hardware resources. This option appears only if the managed system is capable of hardware discovery and the option is selected by default.

   Note: If you do not select the **Retrieve inactive and unallocated hardware resources** option, the HMC does not perform hardware discovery. The HMC still performs inventory gathering and retrieves hardware information for any active partitions on the managed server. The resulting new system plan contains hardware information from the inventory-gathering process, as well as hardware information from the hardware inventory cache on the system.
6. Optional: Select whether you want to view the system plan immediately after the HMC creates it.
7. Click **Create**.

Now that you have a new system plan, you can export the system plan, import it onto another managed system, and deploy the system plan to the managed system.

Note: As an alternative to the HMC Web user interface, you can use the mksysplan command on the HMC to create a system plan based upon the configuration of an existing managed system.

Related tasks

Exporting a system plan using HMC version 7

Importing a system plan using HMC version 7

Deploying a system plan using HMC version 7

Related reference

HMC mksysplan Command

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.

The Integrated Virtualization Manager 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 from the Integrated Virtualization Manager, complete the following steps:

1. From the navigation area, select **Manage System Plans**. The Managed System Plans page opens.
2. Click **Create/Import system plan** from 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.

Now that you have a new system plan, you can export the system plan, import it onto another managed system, and deploy the system plan to the 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.

Related tasks

Exporting a system plan using the Integrated Virtualization Manager

Importing a system plan using the Integrated Virtualization Manager

Deploying a system plan using the Integrated Virtualization Manager

Related reference

IVM mksysplan Command

Specifications

This topic defines the range of configuration possibilities, including the minimum number of resources needed and the maximum number or 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 8. Resources that are required

Resource	Requirement
Hardware Management Console or Integrated Virtualization Manager	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.

Table 9. 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.

Logical volumes exported as Virtual SCSI disks are created by using the Virtual I/O Server command-line interface. If a logical volume is exported as a virtual device, all physical volumes that make up the volume group in which the logical volume is contained need to be attached to the same adapter. You can ensure this situation by creating volume groups with no more than one physical disk.

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.

- Virtual SCSI does not have any limitations in terms of the number of supported adapters. A maximum of 64,000 virtual slots can be assigned to a single partition. Every virtual slot that is created requires resources in order to be instantiated. Therefore, the size of the Virtual I/O Server places a limit on the number of virtual adapters that can be configured.
- 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 using Virtual SCSI devices. Because of the overhead associated with the client/server model, the use of Virtual SCSI can consume additional processor cycles when processing I/O requests.
- The Virtual I/O Server is a dedicated partition, to be used only for Virtual I/O Server operations. Other applications cannot run in the Virtual I/O Server partition.
- If there is a resource shortage, performance degradation might occur. If a Virtual I/O Server is serving many resources to other partitions, ensure that enough processor power is available. In case of high workload across virtual Ethernet adapters and virtual disks, partitions might experience delays in accessing resources.
- Logical volumes exported as Virtual SCSI disks are always configured as single path devices on the client partition.
- Logical volumes exported as Virtual SCSI disks that are part of the root volume group (rootvg) are not persistent when the Virtual I/O Server is updated for maintenance. Therefore, before performing an update procedure, ensure that the corresponding clients' virtual disks are backed up. When exporting logical volumes, it is best to export logical volumes from a volume group other then 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 partitions running only the following operating systems:

- AIX 5.3 or later
- SUSE Linux Enterprise Server 9 for POWER (or later)
- Red Hat Enterprise Linux AS for POWER Version 3 (update 2 or later)
- Red Hat Enterprise Linux AS for POWER Version 4 (or later)

Capacity planning

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

Client 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 partitions, plan what resources each 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 I/O Server resources. However, Virtual SCSI devices generally have higher processor utilization when compared with directly attached storage.

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 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 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.

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.

Related concepts

"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.

"Processor allocation" on page 39
This section contains processor-allocation guidelines for both dedicated processor partitions and shared processor partitions.

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 10. 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)	1000 MB/second

Table 11. 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)	1500 MB/second

Note: For placement rules and limitations, see PCI adapter placement for System p system units and expansion units.

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, 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 physical media overheads, such as inter-frame gaps, preamble bits, cell overhead, data link headers, and trailers. These overheads 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 12. 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
1000 Mb Ethernet, Full Duplex, Host Ethernet Adapter, MTU 1500	10000	9479	1130
1000 Mb Ethernet, Full Duplex, Host Ethernet Adapter, 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 13. 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

Table 13. Two direction (duplex) TCP streaming rates (continued)

Network type	Raw bit rate (Mb)	Payload rate (Mb)	Payload rate (MB)
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, Full Duplex, MTU 1500	10000	14680 (15099 peak)	1750 (1800 peak)
10000 Mb Ethernet, Host Ethernet Adapter, 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.

Related concepts

"Shared Ethernet Adapters" on page 13
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.

Related reference

PCI adapter placement for System p system units and expansion units

Processor allocation:

This section contains processor-allocation guidelines for both dedicated processor partitions and shared processor 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 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 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 partition or a 1.5-processor shared-processor partition.

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

Table 14. 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 15. 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 16. 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 17. 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 overhead. It is approximately 16% to 20% more overhead for MTU 1500 streaming and 31% to 38% more overhead for MTU 9000. The threading option has more overhead 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 partition by itself (without Virtual SCSI in the same 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 partition

Creating a shared-processor 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 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 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 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 partition, add additional processor entitlement as a sizing contingency.

Related reference

mkvdev Command

Memory allocation:

Find information about memory allocation and sizing.

In general, 512 MB of memory per 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.

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 overhead that does not exist with directly attached storage. The overhead 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 overhead 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 partition can be higher and more variable than using a Virtual I/O Server in a dedicated partition. For additional information about the performance differences between dedicated partitions and shared-processor partitions, see Virtual SCSI sizing considerations.

The following table identifies latency overheads for different block-size transmissions on both physical disk and logical-volume-backed Virtual SCSI disks.

Table 18. 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.

Related concepts

"Virtual SCSI sizing considerations" on page 43
Understand the processor and memory-sizing considerations when implementing Virtual SCSI .

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 19. 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 partition or as a dedicated processor 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 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 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 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 multi-threading (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 20. Approximate cycles per second on a 1.65 Ghz 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 partitions on physical disk-backed storage. The first client partition requires a maximum of 7,000 8-KB operations per second. The second client partition requires a maximum of 10,000 8-KB operations per second. The third client 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 partition.

If the I/O rates of the client 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 × 120,000 cycles per operation, resulting in a requirement for approximately 0.19 processor's 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 partitions

Defining Virtual SCSI servers in shared processor partitions allows more specific processor resource sizing and potential recovery of unused processor time by uncapped partitions. However, using shared-processor 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 partition I/O servers, with added entitlement for running in shared-processor 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 partition has high I/O requirements:

- Configure the partition with physical I/O if the configuration allows.
- In most cases, the Virtual I/O Server 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.

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 (also called EtherChannel) 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.

This section contains information about redundancy for both the client 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 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 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 partition to provide highly available access to virtual resources hosted in separate Virtual I/O Server 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 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 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.

When setting up a Virtual SCSI device MPIO configuration on the client logical partition, you must consider the reservation policy of the device on the Virtual I/O Server. To enable an MPIO configuration at the client, none of the Virtual SCSI devices on the Virtual I/O Server should be reserving the SCSI reserve. 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

"Identifying exportable disks" on page 111
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).

"Scenario: Configuring Multi-Path I/O for AIX client logical partitions" on page 89
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.

Related reference

chdev Command

lsdev Command

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 supports certain configurations that utilize the Virtual I/O Server, virtual SCSI and virtual networking capabilities. For the most recent support and configuration information, see the HACMP for System p Web site. For HACMP documentation, see High Availability Cluster Multi-Processing for AIX in the System p5 servers library.

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.

Related concepts

"Link Aggregation or EtherChannel devices" on page 16
A Link Aggregation, or EtherChannel, device is a network port-aggregation technology that allows several Ethernet adapters to be aggregated, which enables them to act as a single Ethernet device. It helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

Related information

HACMP for System p

High Availability Cluster Multi-Processing for AIX

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, which enables them to act as a single Ethernet device. It 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 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.

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 concepts

"Shared Ethernet Adapters" on page 13
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.

Related tasks

"Scenario: Configuring Shared Ethernet Adapter failover" on page 84
Use this article to help you become familiar with typical Shared Ethernet Adapter failover scenario.

Virtual I/O Server 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.

For information about supported storage and multipathing software solutions, see the datasheet available on the Virtual I/O Server Support for UNIX servers and Midrange servers Web site.

Related information

Virtual I/O Server Support for UNIX servers and Midrange servers

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.

See the Virtual I/O Server datasheet see the datasheet available on the Virtual I/O Server Support for UNIX servers and Midrange servers Web site for supported hardware RAID solutions.

Related information

Virtual I/O Server Support for UNIX servers and Midrange servers

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, which enables them to act as a single Ethernet device. It 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.

IBM 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 partitions must be consciously enabled through administrative system-configuration choices.

Using Virtual SCSI on the Virtual I/O Server enables the Virtual I/O Server to provide storage to client 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 partitions can access data on Virtual I/O Server storage devices. For example, a client partition that has access to a logical volume `lv001` exported by the Virtual I/O Server partition cannot access `lv002`, even if it is in the same volume group.

Similar to Virtual SCSI, the firmware also provides the connection between 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 partitions cannot access restricted data.

Related tasks

"Securing the Virtual I/O Server" on page 92
Understand the concepts for securing your Virtual I/O Server environment.

Installing the Virtual I/O Server

Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

These instructions apply to installing the Virtual I/O Server 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 Virtual I/O Server and enabling the Integrated Virtualization Manager.

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. 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.
- Whether you plan to deploy a system plan that includes the Virtual I/O Server. When you deploy a system plan, the HMC automatically creates the Virtual I/O Server logical partition and partition profile and installs the Virtual I/O Server based on the configuration specifications in the system plan.

Related tasks

Determining your HMC machine code version and release

Installing the Virtual I/O Server and enabling the Integrated Virtualization Manager

Related information

System Planning Tool

Installing the Virtual I/O Server by deploying a system plan

When you deploy a system plan that includes the Virtual I/O Server, 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.

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

- 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 partition profile and install the Virtual I/O Server. For instructions, see Installing the Virtual I/O Server manually using the HMC version 6.
- The following items are not configured on the managed system:
  - Client logical partitions
  - Virtual SCSI adapters
  - Virtual Ethernet adapters
  - Shared Ethernet Adapters
  - EtherChannel adapters, or Link Aggregation devices
  - Storage pools
  - Backing devices

  System plans 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.)

  Because of this limitation, the Deploy System Plan wizard will not allow you to partially deploy the Virtual I/O Server logical partition. However, you will still be able to partially deploy other parts of the system plan.

Related tasks

"Installing the Virtual I/O Server manually using the HMC version 6" on page 70
You can create the Virtual I/O Server logical partition and partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 6 or earlier.

Entering the activation code for Advanced POWER Virtualization using the HMC version 7

Use these instructions to enter the activation code using the Hardware Management Console (HMC) version 7 or later.

If your system did not come with the Advanced POWER Virtualization feature enabled, you must use the HMC to enter the activation code that you received when you ordered the feature. This activation code also enables Micro-Partitioning™ on the system.

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 enable Advanced POWER Virtualization. For example, the system on which you plan to install the Virtual I/O Server or enable Micro-Partitioning.
4. Click **Tasks** and select **Capacity on Demand (CoD)** → **Advanced POWER Virtualization** → **Enter Activation Code**.
5. Enter your activation code and click **OK**.

After you enter the activation code, you are ready to create the Virtual I/O Server logical partition and install the Virtual I/O Server. For instructions, see one of the following procedures:

- Deploying a system plan using HMC version 7
- Creating the Virtual I/O Server logical partition and partition profile using the HMC version 7

Related tasks

"Deploying a system plan by using HMC Version 7" on page 54
You can use the Hardware Management Console (HMC) to deploy all or part of a system plan to a managed system.

"Creating the Virtual I/O Server logical partition and partition profile using HMC version 7" on page 61
You can use the Hardware Management Console (HMC) version 7 to create a logical partition and partition profile for the Virtual I/O Server.

Importing a system plan by using 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, such as optical discs or USB drivers, that is mounted on the HMC.
- 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.

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** 2. Click **Import** to display the Upload File window 3. Click **Browse**. 4. Select the system-plan file that you want to import and click **Open**. 5. Click **OK** to upload the file.

Source of the system plan to import	Complete the following steps:
Media	1. Select **Import from media**. 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. 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. 4. Click **Import** to display the Select Media Device window. 5. Select the media that contains the system-plan file you want to import. 6. Click **OK**.
Remote FTP site	1. Select **Import from a remote FTP site**. 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. 3. In the **Remote site hostname** field, enter the host name or IP address of the remote FTP site. 4. In the **User ID** field, enter the user ID to use to access the remote FTP site. 5. In the **Password** field, enter the password to use to access the remote FTP site. 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.

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 using HMC version 7. If you imported the system-plan file from media, you can unmount the media by using the umount command in the HMC command line interface.

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.

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 reference

Tasks and roles

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 partitions in the plan and, if the plan contains Virtual I/O Server provisioning information for a partition, such as storage assignments, the HMC can make these resource assignments for the partition.

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

- 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.
- Verify that the physical hardware and any expansion units are connected and 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.
- 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. 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*.
- If the system plan includes a Storage Area Network (SAN) or Fibre Channel adapters, ensure that the adapters are cabled and the SAN is configured.
- If you plan to deploy the Virtual I/O Server, then ensure that its installation image is on the HMC. To see the installation images on the HMC, enter this command `OS_install -l` on the HMC command line. If the Virtual I/O Server installation image is not listed, then complete the following steps to copy an installation image to the HMC:
  1. Obtain a copy of the Virtual I/O Server on DVD. You can use the original installation media or you can contact your sales representative to obtain another copy. If you cannot obtain a copy of the Virtual I/O Server, you can deploy the remainder of the system plan and install the Virtual I/O Server at a later time.
  2. Insert the DVD into the DVD drive on the HMC.
  3. From the HMC command line, use the OS_install command to copy the Virtual I/O Server installation image from the DVD to the HMC. For example, you can enter the following command:
     ```
     OS_install -o define_resource -a type=AIX -a version=1.4.0.0 -a location=/export/
     resources/vios -a source=/dev/cdrom vios1_install_res.
     ```
- Excluding the Virtual I/O Server logical partitions, shut down any logical partitions that you have already deployed to the managed system from the system plan. For Virtual I/O Server partitions previously deployed, ensure that they are active, and that there is an Resource Monitoring and Control (RMC) connection between the HMC and each Virtual I/O Server partition.
- 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.
- Ensure that you are a super administrator. For information about user roles, refer to Tasks and roles.

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. Click **Tasks** and select **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 does not complete successfully, correct the issues that the error messages describe, click **Cancel** to exit the wizard, and restart this procedure from the beginning.
   d. If the validation process fails, you might want to create a system plan that is based on the current configuration of the managed system. Such a system plan would allow 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.
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 are operating environments 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 partition. At this time, you can select to deploy the Virtual I/O Server operating environment only.
   b. Enter the location of the Virtual I/O Server installation image.
   c. Enter or change late-binding installation settings for the Virtual I/O Server. 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 partition on which you are installing the operating environment.
   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, complete the following tasks:

1. 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. 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.
2. Install operating systems and software on the logical partitions.
3. Configure the virtual I/O adapters that are assigned to each logical partition within the operating systems so that virtual storage resources can be shared among logical partitions.

Related tasks

"Importing a system plan by using HMC Version 7" on page 52
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.

Deleting a logical partition using HMC version 7

Related reference

Tasks and roles

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.

To finish the installation, complete the following steps:

1. Accept the Virtual I/O Server license:
   - If you installed the Virtual I/O Server using the Deploy System Plan wizard on the Hardware Management Console (HMC), then you already accepted the license. Go to the next step.
   - If you did not deploy a system plan, then you must accept the license before you can work with the Virtual I/O Server. For instructions, see Viewing and accepting the Virtual I/O Server license.
2. Check for updates to the Virtual I/O Server. For instructions, see Updating the Virtual I/O Server.
3. Set up remote connections to the Virtual I/O Server. For instructions, see Connecting to the Virtual I/O Server using OpenSSH.
4. 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.
5. 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.

When you are finished, you are ready to configure virtual SCSI and shared Ethernet resources. For instructions, see Managing the Virtual I/O Server.

Related tasks

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

"Viewing and accepting the Virtual I/O Server license" on page 66
You must view and accept the license before using the Virtual I/O Server.

"Updating the Virtual I/O Server" on page 127
Find instructions for updating the Virtual I/O Server.

"Connecting to the Virtual I/O Server using OpenSSH"
You can set up remote connections to the Virtual I/O Server using secure connections.

"Managing users on the Virtual I/O Server" on page 118
Find commands for creating, listing, changing, switching, and removing users.

"Managing the Virtual I/O Server" on page 95
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.

Related reference

mktcpip Command

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.

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.
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.

Related tasks

"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.

Related information

OpenSSL Project

Portable OpenSSH

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.

Related information

OpenSSL Project

Portable OpenSSH

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. To download the OpenSSH software, complete 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`.

Related reference

mkdir Command

startnetsvc Command

Related information

AIX Toolbox for Linux Applications

SourceFORGE.net

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.

Related reference

updateios Command

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

You can create the Virtual I/O Server logical partition and 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.

Related tasks

"Installing the Virtual I/O Server by deploying a system plan" on page 51
When you deploy a system plan that includes the Virtual I/O Server, 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.

"Installing the Virtual I/O Server manually using the HMC version 6" on page 70
You can create the Virtual I/O Server logical partition and partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 6 or earlier.

Entering the activation code for Advanced POWER Virtualization using the HMC version 7

Use these instructions to enter the activation code using the Hardware Management Console (HMC) version 7 or later.

If your system did not come with the Advanced POWER Virtualization feature enabled, you must use the HMC to enter the activation code that you received when you ordered the feature. This activation code also enables Micro-Partitioning on the system.

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 enable Advanced POWER Virtualization. For example, the system on which you plan to install the Virtual I/O Server or enable Micro-Partitioning.
4. Click **Tasks** and select **Capacity on Demand (CoD)** → **Advanced POWER Virtualization** → **Enter Activation Code**.
5. Enter your activation code and click **OK**.

After you enter the activation code, you are ready to create the Virtual I/O Server logical partition and install the Virtual I/O Server. For instructions, see one of the following procedures:

- Deploying a system plan using HMC version 7
- Creating the Virtual I/O Server logical partition and partition profile using the HMC version 7

Related tasks

"Deploying a system plan by using HMC Version 7" on page 54
You can use the Hardware Management Console (HMC) to deploy all or part of a system plan to a managed system.

"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.

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 Advanced POWER Virtualization feature is activated. For instructions, see Entering the activation code for Advanced POWER Virtualization using the HMC version 7.

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.
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

Related concepts

Logical partition profile

"Host Ethernet Adapter" on page 14
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).

Related tasks

"Entering the activation code for Advanced POWER Virtualization using the HMC version 7" on page 51
Use these instructions to enter the activation code using the Hardware Management Console (HMC) version 7 or later.

"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.

"Installing the Virtual I/O Server from CD or DVD" on page 63
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.

Related reference

Tasks and roles

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

Related tasks

“Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6” on page 70
You can use the Hardware Management Console (HMC) version 6 to create a logical partition and partition profile for the Virtual I/O Server.

“Creating the Virtual I/O Server logical partition and partition profile using HMC version 7” on page 61
You can use the Hardware Management Console (HMC) version 7 to create a logical partition and partition profile for the Virtual I/O Server.

“Finishing the Virtual I/O Server installation” on page 56
After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

“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.

Related reference

Tasks and roles

installios Command

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 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
  - Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6
- 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 resides.
     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 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 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 partition to open the menu.
    c. Click **Activate**. The Activate Partition menu opens with a selection of 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 partition.

2. Activate the Virtual I/O Server logical partition:
   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 partition to open the menu.
   c. Click **Activate**. The Activate Partition menu opens with a selection of 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 partition.
3. 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.
4. 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.

Related tasks

"Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6" on page 70
You can use the Hardware Management Console (HMC) version 6 to create a logical partition and partition profile for the Virtual I/O Server.

"Creating the Virtual I/O Server logical partition and partition profile using HMC version 7" on page 61
You can use the Hardware Management Console (HMC) version 7 to create a logical partition and partition profile for the Virtual I/O Server.

"Finishing the Virtual I/O Server installation" on page 56
After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

"Installing the Virtual I/O Server from the HMC" on page 62
Find instructions for installing the Virtual I/O Server from the HMC by using the **installios** command.

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.

To finish the installation, complete the following steps:

1. Accept the Virtual I/O Server license:
   - If you installed the Virtual I/O Server using the Deploy System Plan wizard on the Hardware Management Console (HMC), then you already accepted the license. Go to the next step.
   - If you did not deploy a system plan, then you must accept the license before you can work with the Virtual I/O Server. For instructions, see Viewing and accepting the Virtual I/O Server license.
2. Check for updates to the Virtual I/O Server. For instructions, see Updating the Virtual I/O Server.
3. Set up remote connections to the Virtual I/O Server. For instructions, see Connecting to the Virtual I/O Server using OpenSSH.
4. 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.
5. 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.

When you are finished, you are ready to configure virtual SCSI and shared Ethernet resources. For instructions, see Managing the Virtual I/O Server.

Related tasks

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

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

"Updating the Virtual I/O Server" on page 127
Find instructions for updating the Virtual I/O Server.

"Connecting to the Virtual I/O Server using OpenSSH" on page 57
You can set up remote connections to the Virtual I/O Server using secure connections.

"Managing users on the Virtual I/O Server" on page 118
Find commands for creating, listing, changing, switching, and removing users.

"Managing the Virtual I/O Server" on page 95
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.

Related reference

mktcpip Command

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 partition profile is created and the Virtual I/O Server is installed. For instructions, see Installing the Virtual I/O Server.

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

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.
3. In the installation program, English is the default language. If you need to change the language setting for the system, follow these steps. Otherwise, proceed to step 4.
   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
      ```
4. View the license by typing `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:
   a. View the list of available locales to display the license by typing the following command:
      ```
      license -ls
      ```
   b. 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
```

5. Accept the license agreement for the Virtual I/O Server by typing the following command:

   ```
   license -accept -lang Name
   ```

Related tasks

"Finishing the Virtual I/O Server installation" on page 56
After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

Related reference

chlang Command

license Command

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.

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.
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.

Related tasks

"Downloading, installing, and updating OpenSSH and OpenSSL" on page 58
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.

Related information

OpenSSL Project

Portable OpenSSH

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.

Related information

OpenSSL Project

Portable OpenSSH

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. To download the OpenSSH software, complete 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`.

Related reference

mkdir Command

startnetsvc Command

Related information

AIX Toolbox for Linux Applications

SourceFORGE.net

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.

Related reference

updateios Command

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

You can create the Virtual I/O Server logical partition and 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 by deploying a system plan
  – Installing the Virtual I/O Server manually using the HMC version 7

Related tasks

"Installing the Virtual I/O Server by deploying a system plan" on page 51
When you deploy a system plan that includes the Virtual I/O Server, 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.

"Installing the Virtual I/O Server manually using the HMC version 7" on page 60
You can create the Virtual I/O Server logical partition and partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 7 or later.

Entering the activation code for Advanced POWER Virtualization using the HMC version 6

The 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 your system did not come with the Advanced POWER Virtualization feature enabled, you must use the HMC to enter the activation code that you received when you ordered the feature. This activation code also enables Micro-Partitioning 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.

After you enter the activation code, you are ready to create the Virtual I/O Server logical partition and install the Virtual I/O Server. For instructions, see Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6.

Related tasks

"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.

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 Advanced POWER Virtualization feature is activated. For instructions, see Entering the activation code for Advanced POWER Virtualization using the HMC version 6.

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

Related concepts

Logical partition profile

Related tasks

"Entering the activation code for Advanced POWER Virtualization using the HMC version 6" on page 70
The 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).

"Installing the Virtual I/O Server from the HMC" on page 62
Find instructions for installing the Virtual I/O Server from the HMC by using the **installios** command.

"Installing the Virtual I/O Server from CD or DVD" on page 63
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.

Related reference

Tasks and roles

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

Related tasks

"Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6" on page 70
You can use the Hardware Management Console (HMC) version 6 to create a logical partition and partition profile for the Virtual I/O Server.

"Creating the Virtual I/O Server logical partition and partition profile using HMC version 7" on page 61
You can use the Hardware Management Console (HMC) version 7 to create a logical partition and partition profile for the Virtual I/O Server.

"Finishing the Virtual I/O Server installation" on page 56
After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

"Installing the Virtual I/O Server from CD or DVD" on page 63
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.

Related reference

Tasks and roles

installios Command

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 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
  - Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6
- 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 resides.
     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 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 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 partition to open the menu.

   c. Click **Activate**. The Activate Partition menu opens with a selection of 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 partition.

2. Activate the Virtual I/O Server logical partition:
   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 partition to open the menu.
   c. Click **Activate**. The Activate Partition menu opens with a selection of 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 partition.
3. 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.
4. 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.

Related tasks

"Creating the Virtual I/O Server logical partition and partition profile using the HMC version 6" on page 70
You can use the Hardware Management Console (HMC) version 6 to create a logical partition and partition profile for the Virtual I/O Server.

"Creating the Virtual I/O Server logical partition and partition profile using HMC version 7" on page 61
You can use the Hardware Management Console (HMC) version 7 to create a logical partition and partition profile for the Virtual I/O Server.

"Finishing the Virtual I/O Server installation" on page 56
After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

"Installing the Virtual I/O Server from the HMC" on page 62
Find instructions for installing the Virtual I/O Server from the HMC by using the **installios** command.

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.

To finish the installation, complete the following steps:

1. Accept the Virtual I/O Server license:
   - If you installed the Virtual I/O Server using the Deploy System Plan wizard on the Hardware Management Console (HMC), then you already accepted the license. Go to the next step.
   - If you did not deploy a system plan, then you must accept the license before you can work with the Virtual I/O Server. For instructions, see Viewing and accepting the Virtual I/O Server license.
2. Check for updates to the Virtual I/O Server. For instructions, see Updating the Virtual I/O Server.
3. Set up remote connections to the Virtual I/O Server. For instructions, see Connecting to the Virtual I/O Server using OpenSSH.
4. 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.
5. 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.

When you are finished, you are ready to configure virtual SCSI and shared Ethernet resources. For instructions, see Managing the Virtual I/O Server.

Related tasks

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

"Viewing and accepting the Virtual I/O Server license" on page 66
You must view and accept the license before using the Virtual I/O Server.

"Updating the Virtual I/O Server" on page 127
Find instructions for updating the Virtual I/O Server.

"Connecting to the Virtual I/O Server using OpenSSH" on page 57
You can set up remote connections to the Virtual I/O Server using secure connections.

"Managing users on the Virtual I/O Server" on page 118
Find commands for creating, listing, changing, switching, and removing users.

"Managing the Virtual I/O Server" on page 95
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.

Related reference

mktcpip Command

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 partition profile is created and the Virtual I/O Server is installed. For instructions, see Installing the Virtual I/O Server.

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

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.
3. In the installation program, English is the default language. If you need to change the language setting for the system, follow these steps. Otherwise, proceed to step 4.
   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
      ```
4. View the license by typing `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:
   a. View the list of available locales to display the license by typing the following command:
      ```
      license -ls
      ```
   b. 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
```

5. Accept the license agreement for the Virtual I/O Server by typing the following command:

   ```
   license -accept -lang Name
   ```

Related tasks

"Finishing the Virtual I/O Server installation" on page 56
After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

Related reference

chlang Command

license Command

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.

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.
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.

Related tasks

"Downloading, installing, and updating OpenSSH and OpenSSL" on page 58
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.

Related information

OpenSSL Project

Portable OpenSSH

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.

Related information

OpenSSL Project

Portable OpenSSH

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. To download the OpenSSH software, complete 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`.

Related reference

mkdir Command

startnetsvc Command

Related information

AIX Toolbox for Linux Applications

SourceFORGE.net

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.

Related reference

updateios Command

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. For setup instructions, see Installing the HMC.
- You understand the partitioning concepts as described in Concepts for Partitioning the server.
- The Virtual I/O Server partition has been created and the Virtual I/O Server has been installed. For instructions, see Installing the Virtual I/O Server.
- 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.

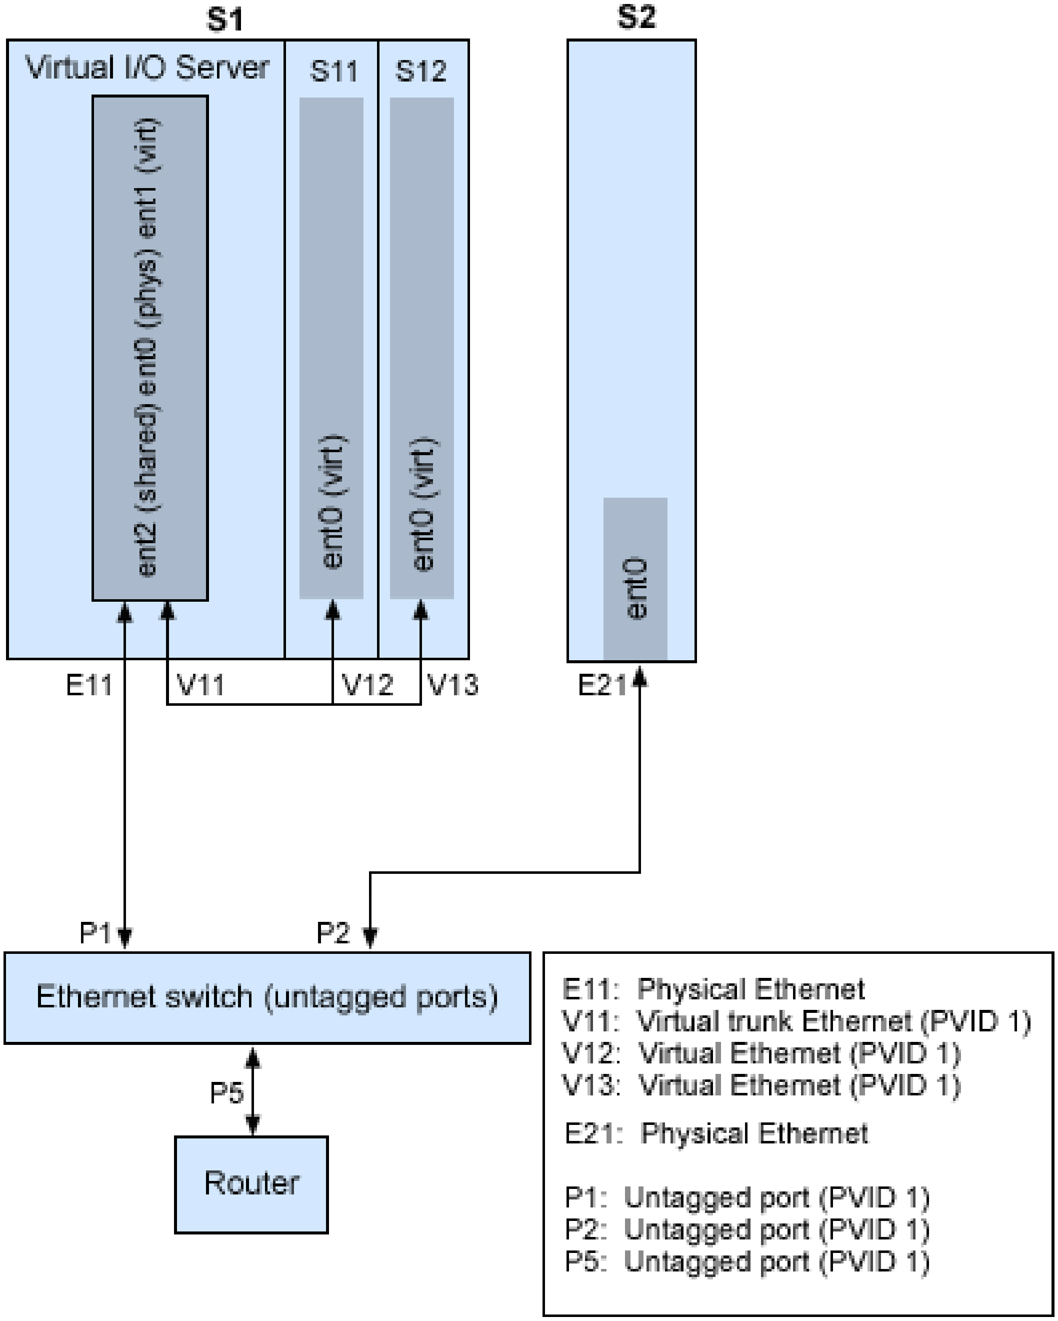
S1
Virtual I/O Server
ent2 (shared) ent0 (phys) ent1 (virt)
S11
ent0 (virt)
S12
ent0 (virt)
S2
ent0
E11
V11
V12
V13
E21
P1
P2
Ethernet switch (untagged ports)
P5
Router
E11: Physical Ethernet
V11: Virtual trunk Ethernet (PVID 1)
V12: Virtual Ethernet (PVID 1)
V13: Virtual Ethernet (PVID 1)
E21: Physical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

Related concepts

Concepts for Partitioning the server

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

Related tasks

Installing the HMC

"Installing the Virtual I/O Server manually using the HMC version 6" on page 70
You can create the Virtual I/O Server logical partition and partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 6 or earlier.

Related information

Managing your server using the HMC

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 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. For setup instructions, see Installing the HMC.
- You understand the partitioning concepts as described in Concepts for Partitioning the server.
- The Virtual I/O Server partition has been created and the Virtual I/O Server has been installed. For instructions, see Installing the Virtual I/O Server.
- 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 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.

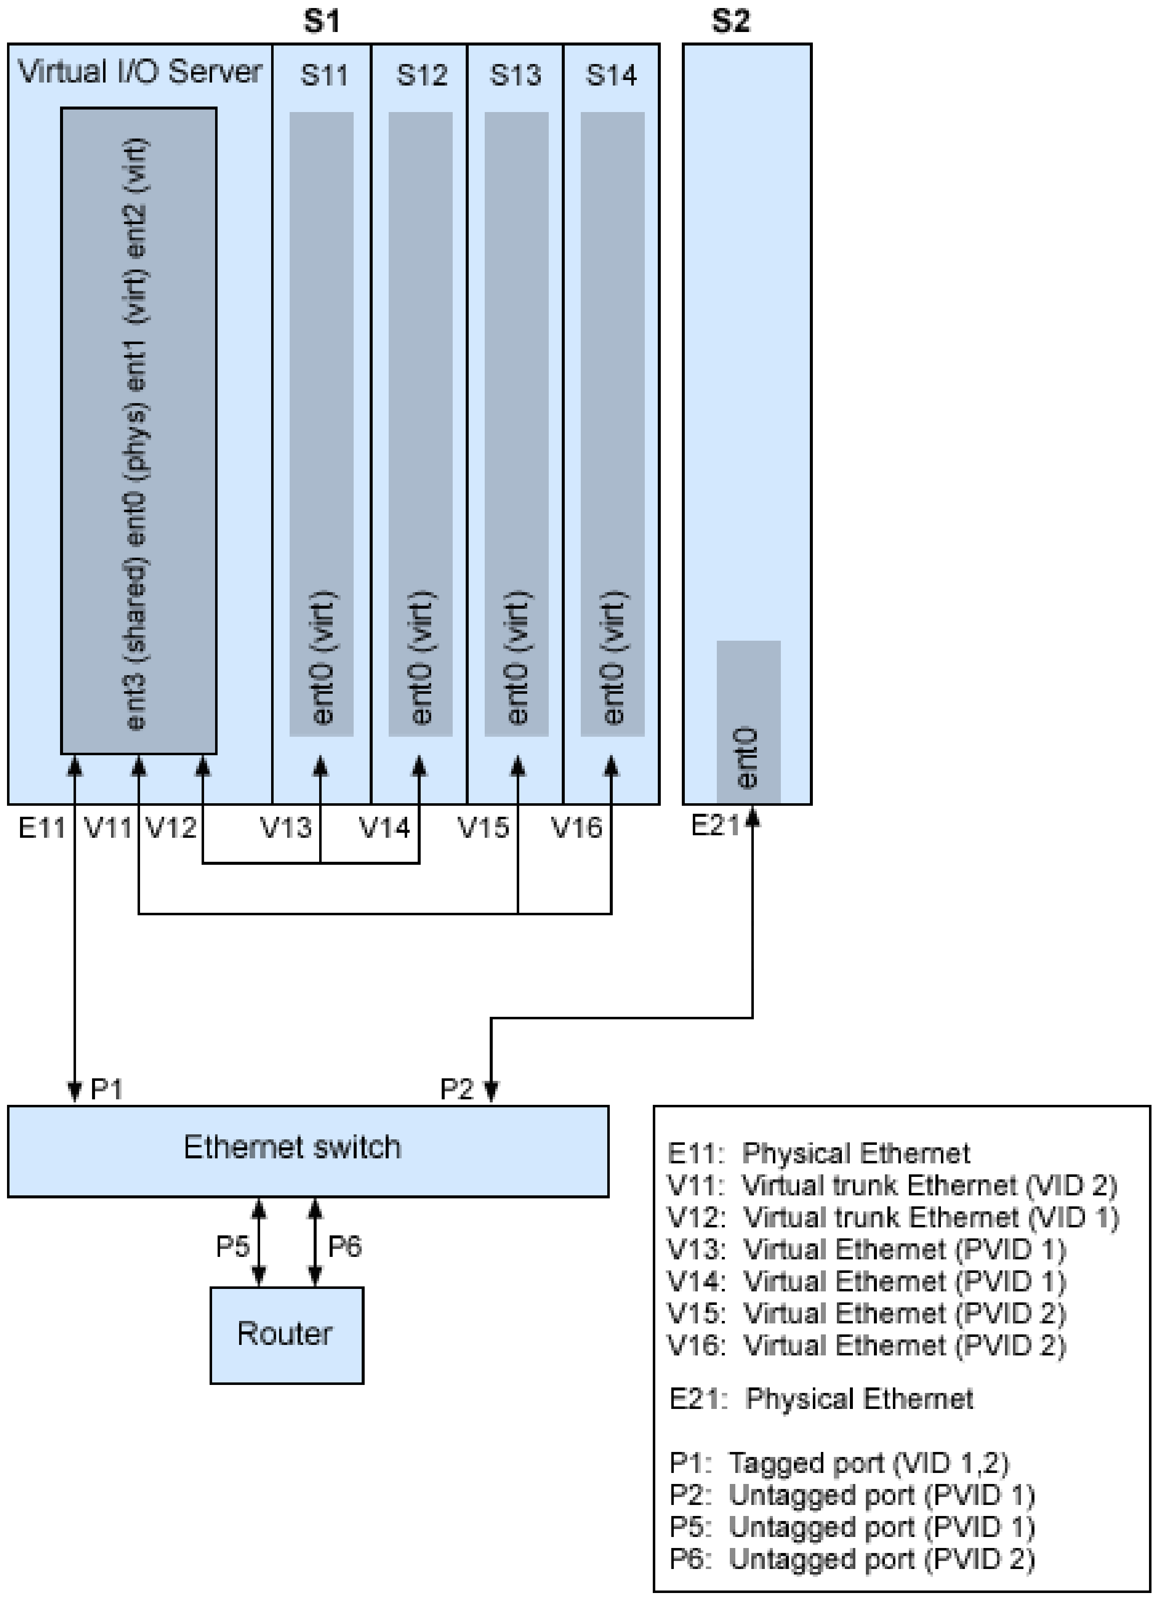
S1
S2
Virtual I/O Server
S11
S12
S13
S14
ent3 (shared) ent0 (phys) ent1 (virt) ent2 (virt)
ent0 (virt)
ent0 (virt)
ent0 (virt)
ent0 (virt)
ent0
E11
V11
V12
V13
V14
V15
V16
E21
P1
P2
Ethernet switch
P5
P6
Router
E11: Physical Ethernet
V11: Virtual trunk Ethernet (VID 2)
V12: Virtual trunk Ethernet (VID 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 partitions:
   - Create virtual adapters V13 and V14 for partitions S11 and S12, respectively, with PVID set to 2 and no additional VIDs.
   - Create virtual adapters V15 and V16 for 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 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.

Related concepts

Concepts for Partitioning the server

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

Related tasks

Installing the HMC

"Installing the Virtual I/O Server manually using the HMC version 6" on page 70
You can create the Virtual I/O Server logical partition and partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 6 or earlier.

Related information

Managing your server using the HMC

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 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 partitions will not be lost in the case of adapter failure.

Prerequisites and assumptions

- The Hardware Management Console (HMC) was set up. For setup instructions, see Installing the HMC.
- You understand the partitioning concepts as described in Concepts for Partitioning the server.
- Two separate Virtual I/O Server partitions have been created and the Virtual I/O Server has been installed in each. For instructions, see Installing the Virtual I/O Server.
- You understand what Shared Ethernet Adapter failover is and how it works. See Shared Ethernet Adapter failover.
- You have created the remaining logical partitions that you want added to the network configuration.
- EachVirtual I/O Server partition has an available physical Ethernet adapter assigned to it.
- You have IP addresses for all partitions and systems that will be added to the configuration.

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

The following image depicts a configuration where the Shared Ethernet Adapter failover feature is set up. The client 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.

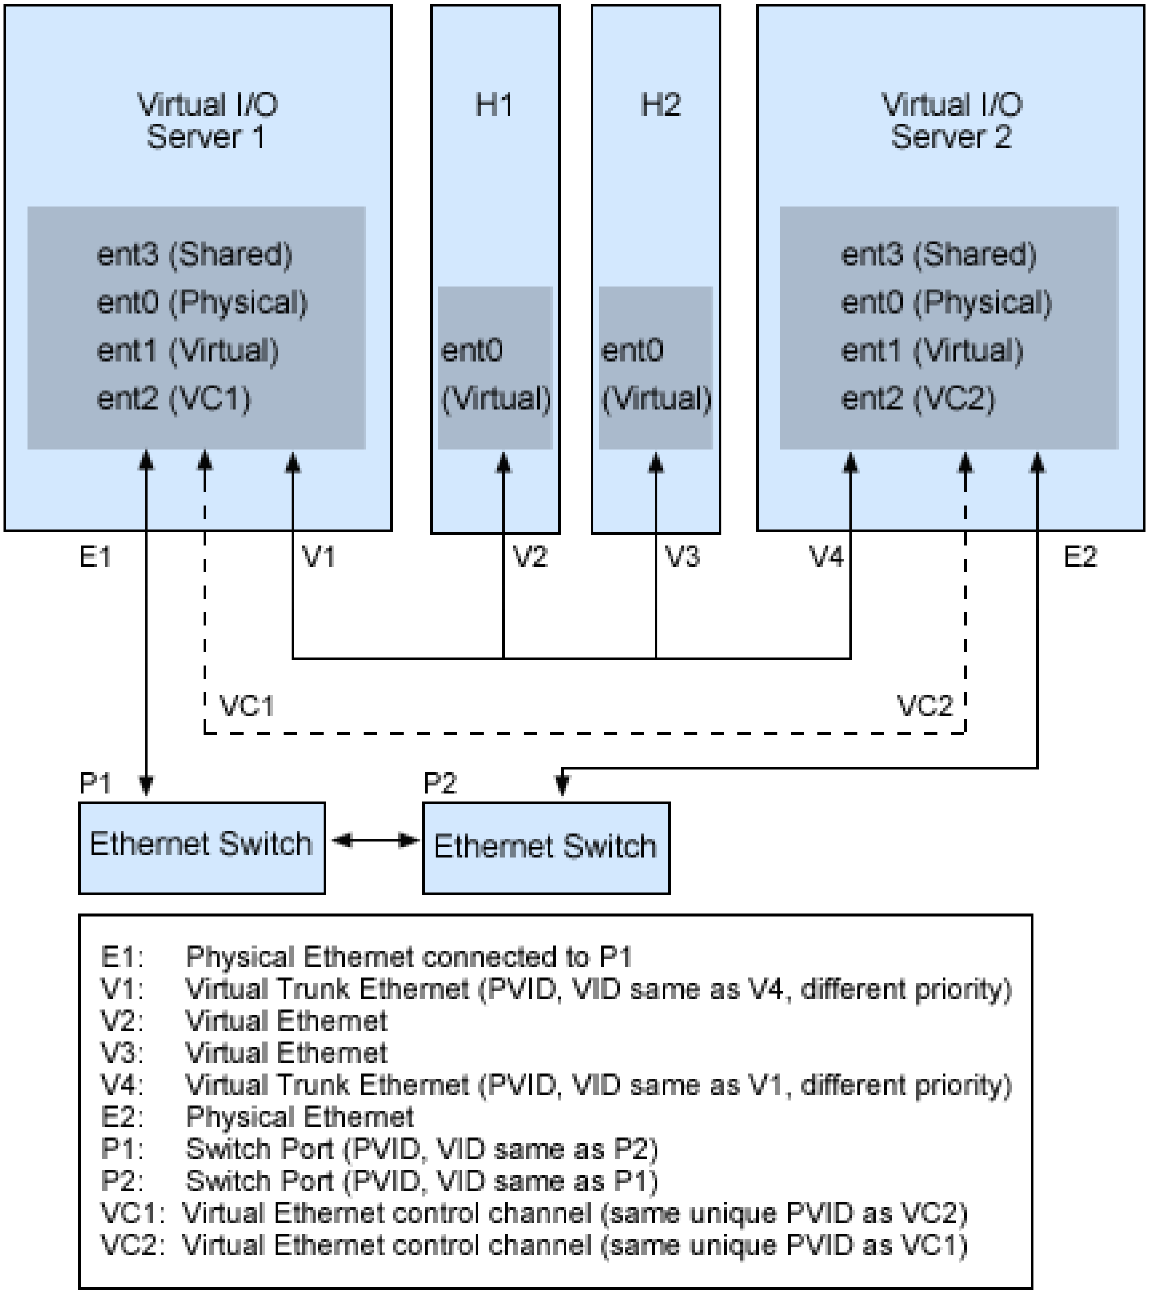
Virtual I/O Server 1
ent3 (Shared)
ent0 (Physical)
ent1 (Virtual)
ent2 (VC1)
H1
ent0 (Virtual)
H2
ent0 (Virtual)
Virtual I/O Server 2
ent3 (Shared)
ent0 (Physical)
ent1 (Virtual)
ent2 (VC2)
E1
V1
V2
V3
V4
E2
VC1
VC2
P1
P2
Ethernet Switch
Ethernet Switch
E1: Physical Ethernet connected to P1
V1: Virtual 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 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 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 partitions:

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

Related concepts

Concepts for Partitioning the server

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

"Shared Ethernet Adapter failover" on page 48
Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption.

Related tasks

Installing the HMC

"Installing the Virtual I/O Server manually using the HMC version 6" on page 70
You can create the Virtual I/O Server logical partition and partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 6 or earlier.

Related reference

mkvdev Command

Related information

Managing your server using the HMC

Scenario: Configuring Network Interface Backup in Virtual I/O clients 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 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 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:

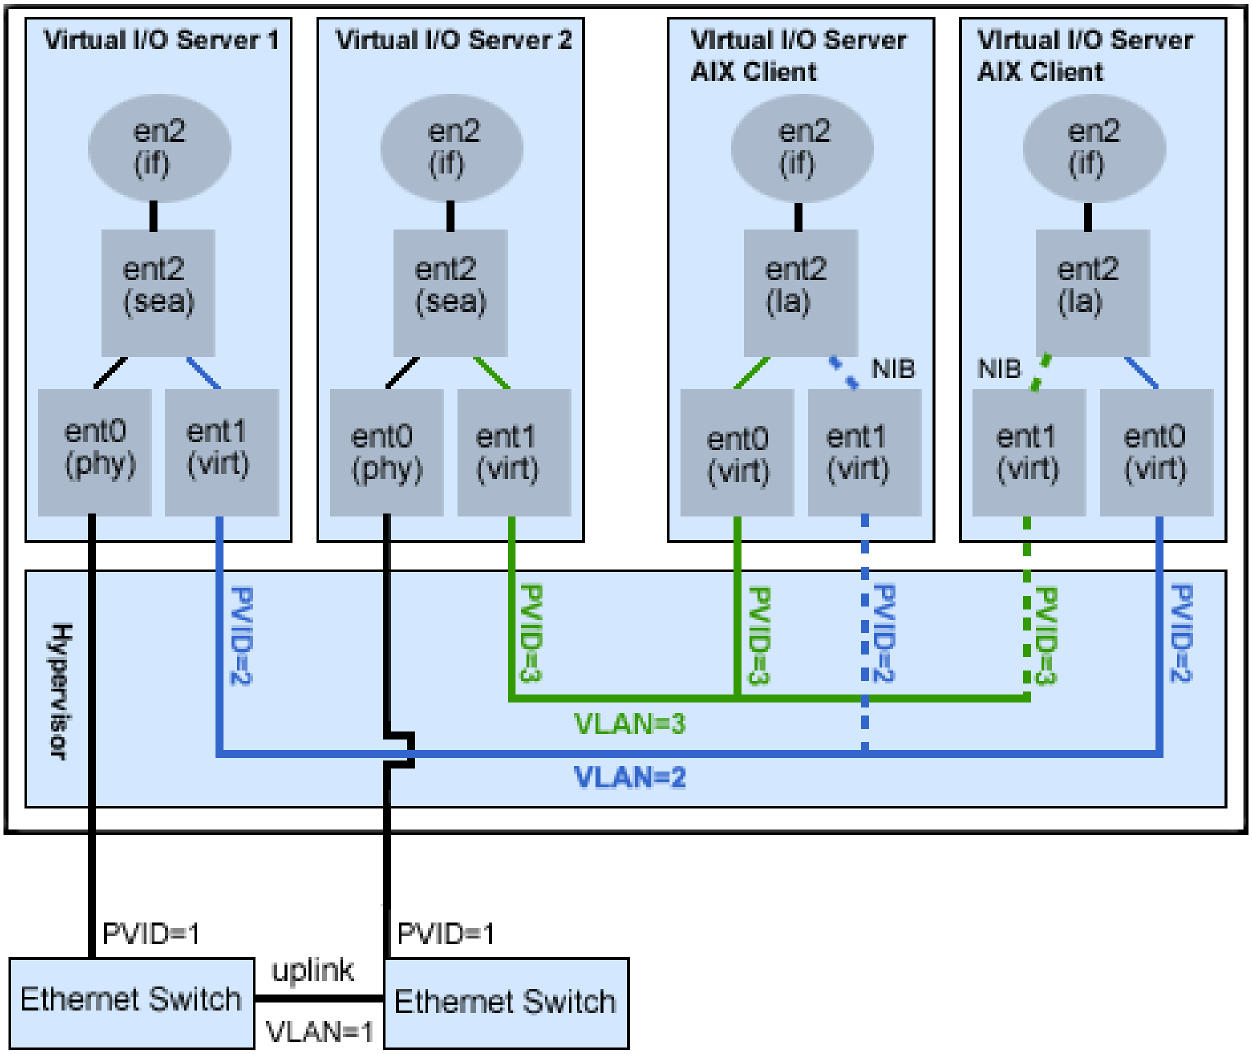

Prerequisites and assumptions

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

- The Hardware Management Console (HMC) is already set up. For setup instructions, see Installing the HMC.
- Two separate Virtual I/O Server partitions have been created and the Virtual I/O Server has been installed in each. See the instructions in Installing the Virtual I/O Server.
- You have created the remaining logical partitions that you want added to the network configuration.
- Each Virtual I/O Server partition has an available physical Ethernet adapter assigned to it.
- You have IP addresses for all 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.
   b. Configure a Shared Ethernet Adapter on the second Virtual I/O Server, as in step 1.
2. For each client 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 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. To create this configuration, follow the procedure Configuring an EtherChannel in the IBM System p and AIX Information Center. 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 configuring NIB with two virtual Ethernet adapters, the internal networks used must stay separated in the POWER Hypervisor, so you must use different PVIDs for the two adapters in the client and cannot use additional VIDs on them.

Related concepts

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

Related tasks

Installing the HMC

"Installing the Virtual I/O Server manually using the HMC version 6" on page 70
You can create the Virtual I/O Server logical partition and partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 6 or earlier.

"Configuring a Shared Ethernet Adapter" on page 97
Find instructions for configuring Shared Ethernet Adapters.

Related information

Managing your server using the HMC

Configuring an EtherChannel

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.

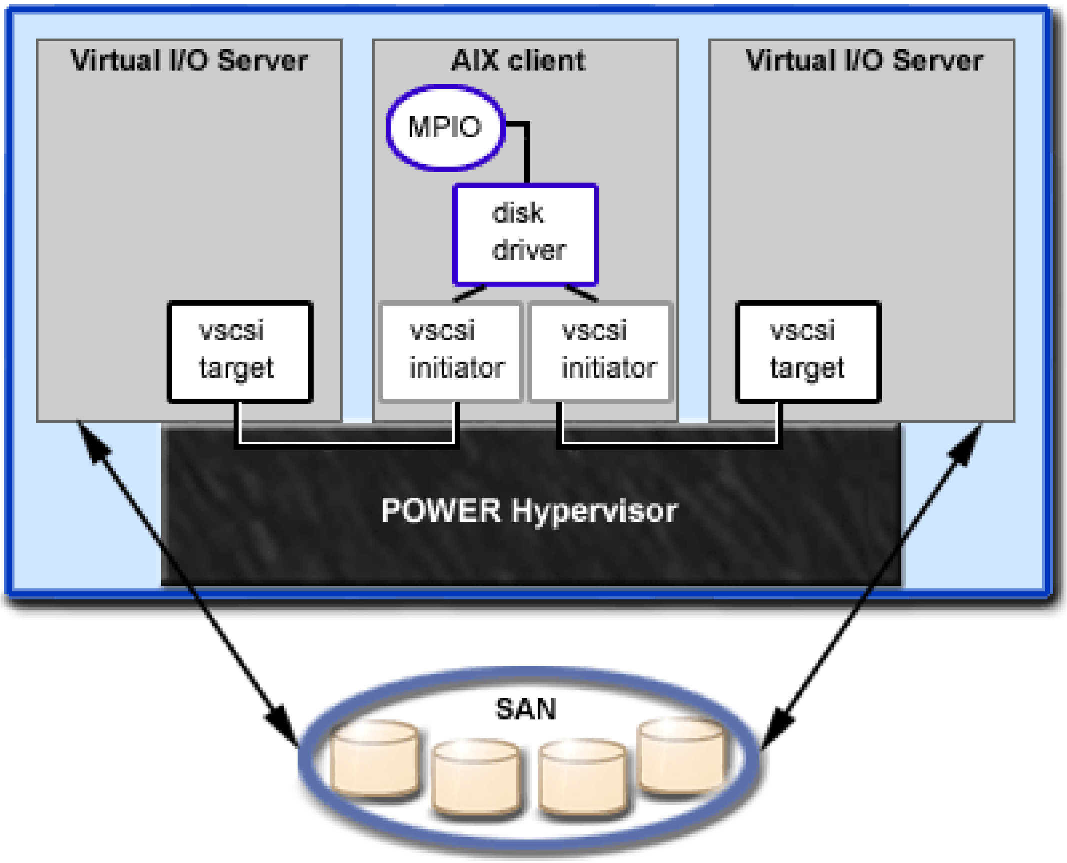

Using the preceding figure as a guide, follow these steps:

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. 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 hdisk`*`x`*` -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 hdisk`*`x`*` -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.

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 hdisk`*`x`*` -a hcheck_interval=60 -P`. The client logical partition must be rebooted for this change to take effect.

Related tasks

"Identifying exportable disks" on page 111
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).

Related reference

chdev Command

lsdev Command

lspath Command

lspv Command

mkvdev Command

Related information

Installing AIX in a Partitioned Environment

Securing the Virtual I/O Server

Understand the concepts for securing your Virtual I/O Server environment.

The Virtual I/O Server (VIOS) provides extra security features that enable you to control access to the virtual environment and ensure the security of your system. These features are available with Virtual I/O Server version 1.3 or later. The following topics discuss the security features available and provide tips for ensuring a secure environment for your Virtual I/O Server setup.

Introduction to Virtual I/O Server security

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.

The viosecure command enables you to set, change, and view current security settings. The settings are not enabled by default, you must run the viosecure command to specify the options.

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.

The system security hardening features provided by Virtual I/O Server enable you to specify values such as the following:

- password policy settings
- usrck, pwdck, grpck, and sysck actions
- Default file creation settings
- crontab settings

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

The Virtual I/O Server firewall enables you to 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.

Related tasks

"Configuring Virtual I/O Server system security hardening"
Set the security level to specify security hardening rules for your Virtual I/O Server system.

"Configuring Virtual I/O Server firewall settings" on page 94
Enable the Virtual I/O Server firewall to control IP activity.

Related reference

man Command

viosecure Command

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 enable the settings.

Use the following tasks to configure the system security settings.

Related reference

viosecure Command

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`

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
```

Related reference

viosecure Command

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.

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 Partitioning with the Integrated Virtualization Manager.

Related tasks

Partitioning with the Integrated Virtualization Manager

Managing virtual Ethernet

This topic contains configuration and management information for virtual Ethernet, including configuring Shared Ethernet Adapters, configuring Link Aggregation (or EtherChannel), and changing network settings.

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, 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.

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 **Properites**. The Logical Partition Profile Properties page is displayed.
5. Click the**Virtual Adapters** tab.
6. Click **Virtual Ethernet**.
7. Click **New**
8. Select **IEEE 802.1Q-compatible adapter**.
9. 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.
10. 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.
11. 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.
12. When you are finished, click **Submit**.
13. 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.

14. Click **OK** to exit the Logical Partition Profile Properties page.
15. Click **Close** to exit the Managed Profiles page.
16. 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.

Related concepts

"Shared Ethernet Adapters" on page 13
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.

Related tasks

"Configuring a Shared Ethernet Adapter" on page 97
Find instructions for configuring Shared Ethernet Adapters.

"Setting the LHEA to promiscuous mode"
To use a Shared Ethernet Adapter with a Host Ethernet Adapter, you must set the Logical Host Ethernet Adapter (LHEA) to promiscuous mode.

Setting the LHEA to promiscuous mode:

To use a Shared Ethernet Adapter with a Host Ethernet Adapter, 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.

Related concepts

"Shared Ethernet Adapters" on page 13
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.

"Host Ethernet Adapter" on page 14
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).

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.

Related concepts

"Shared Ethernet Adapters" on page 13
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.

Related tasks

"Configuring a Shared Ethernet Adapter"
Find instructions for configuring Shared Ethernet Adapters.

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
- Creating a Shared Ethernet Adapter using HMC version 6

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:

- You can also use a Link Aggregation, or EtherChannel, device as the Shared Ethernet Adapter.
- If you plan to use the Host Ethernet Adapter 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* is the physical adapter being used as part of the Shared Ethernet Adapter device.
   - *virtual_ethernet_adapters* are the virtual Ethernet adapter or adapters that will use the Shared Ethernet Adapter.
   - *DefaultVirtualEthernetAdapter* is the default virtual Ethernet adapter used to handle untagged packets. If you have only one virtual Ethernet adapter for this partition, use it as the default.
   - *SEADefaultPVID* is 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 (for example, using Telnet)?
   - 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 partitions' operating systems, those partitions can communicate with the external network.

Related concepts

"Shared Ethernet Adapters" on page 13
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.

Related tasks

"Creating a Shared Ethernet Adapter using HMC version 6" on page 97
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.

"Creating a Shared Ethernet Adapter using HMC version 7" on page 95
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.

"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.

Related reference

chdev Command

lsdev Command

mktcpip Command

mkvdev Command

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.

Related concepts

"Link Aggregation or EtherChannel devices" on page 16
A Link Aggregation, or EtherChannel, device is a network port-aggregation technology that allows several Ethernet adapters to be aggregated, which enables them to act as a single Ethernet device. It helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

Related reference

cfglnagg Command

mkvdev Command

Changing the network configuration

Follow these steps to change or remove the network settings on the Virtual I/O Server partition, such as the IP address, subnet mask, gateway, and nameserver address

In this scenario, the Virtual I/O Server partition already has its network configuration set. The current configuration will be removed, and the updated configuration will then be set.

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.

For example, the Virtual I/O Server partition needs to have its 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.

Related reference

lstcpip Command

mktcpip Command

rmtcpip Command

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.

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.

Related concepts

"Shared Ethernet Adapters" on page 13
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.

Related tasks

"Creating a Shared Ethernet Adapter using HMC version 7" on page 95
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.

"Creating a Shared Ethernet Adapter using HMC version 6" on page 97
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.

Related reference

chdev Command

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, SNMP enables network administrators to more easily manage their networks for the following reasons:

- It hides the underlying system network
- It enables the administrator to 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 21. Tasks and associated commands for working with SNMP on the Virtual I/O Server

Task	Command
Enable SNMP	startnetsvc
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 reference

cl_snmp Command

snmp_info Command

snmp_trap Command

startnetsvc Command

stopnetsvc Command

Related information

Network Management

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. **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*) **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. 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. 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.

Attribute	Description
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.
Thread (*thread*)	Activates or deactivates threading on the Shared Ethernet Adapter. Activating this option adds approximately 16% to 20% more overhead for MTU 1500 streaming and 31% to 38% more overhead for MTU 9000. The threading option has more overhead 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. Threaded mode should be used when Virtual SCSI will be run on the same Virtual I/O Server 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 partition will be dedicated to running shared Ethernet devices (and associated virtual Ethernet devices) only, the adapters should be configured with threading disabled. 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`
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 partition to the Shared Ethernet Adapter to work. Also, the partition must be capable of performing a largesend operation. On AIX, largesend can be enabled on a partition using the ifconfig command. 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. For example, the following command enables *largesend* for Shared Ethernet Adapter ent1: `chdev -l ent1 -a largesend=1` 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 (HEA).

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. 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. If the host uses a domain name server for name resolution, the host name must contain the full domain name. 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: `host.subdomain.subdomain.rootdomain` 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*]. **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. **gateway** Identifies the gateway to which packets are addressed. The *Gateway* parameter can be specified either by symbolic name or numeric address. **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 • 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. 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 • 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. 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 • 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. For the adapter to run at 1000 Mbit/s, the autonegotiation setting must be selected. **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 • Gigabit Ethernet-SX PCI-X Adapter Device Driver • 2-Port 10/100/1000 Base-TX PCI-X Adapter • 2-Port Gigabit Ethernet-SX PCI-X Adapter • 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 • Gigabit Ethernet-SX PCI-X Adapter Device Driver • 2-Port 10/100/1000 Base-TX PCI-X Adapter • 2-Port Gigabit Ethernet-SX PCI-X Adapter • 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 • Gigabit Ethernet-SX PCI-X Adapter Device Driver • 2-Port 10/100/1000 Base-TX PCI-X Adapter • 2-Port Gigabit Ethernet-SX PCI-X Adapter • Gigabit Ethernet-SX PCI Adapter Device Driver • 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. 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. 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: • 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. • 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. 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 • Gigabit Ethernet-SX PCI-X Adapter Device Driver • 2-Port 10/100/1000 Base-TX PCI-X Adapter • 2-Port Gigabit Ethernet-SX PCI-X Adapter • 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. The 8023ad mode enables the Link Aggregation Control Protocol (LACP) to negotiate the adapters in the Link Aggregation device with an LACP-enabled switch. If the Hash Mode attribute is set to anything other than default, this attribute must be set to standard or 8023ad. 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: • `default`: uses the destination IP address to determine the outgoing adapter. • `src_port`: uses the source TCP or UDP port for that connection. • `dst_port`: uses the destination TCP or UDP port for that connection. • `src_dst_port`: uses both the source and destination TCP or UDP ports for that connection to determine the outgoing adapter. 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. If the packet is not TCP or UDP, it uses the default hashing mode (destination IP address). 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. Setting this to yes will enable you to specify a MAC address that you want the Link Aggregation device to use. If you set this option to no, the Link Aggregation device will use 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.

Related concepts

"Shared Ethernet Adapters" on page 13
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.

Related tasks

"Enabling and disabling GVRP" on page 100
You can enable and disable GARP VLAN Registration Protocol (GVRP) on your Shared Ethernet Adapters to control dynamic registration of VLANs over networks.

Related reference

cfglnagg Command

chdev Command

mkvdev Command

Managing virtual SCSI

Find instructions for managing virtual storage devices and logical volumes.

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 partition as a whole or can be partitioned into logical volumes. These logical volumes can be exported as virtual disks to one or more client partitions. Therefore, virtual SCSI enables sharing of adapters and disk devices.

To make a physical or 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 SCSI client adapter is linked to a particular virtual SCSI server adapter in the Virtual I/O Server partition. The client 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 partition.

For information about supported SCSI devices, see the Virtual I/O Server support site at http://techsupport.services.ibm.com/server/virtualization/vios.

Related information

Virtual I/O Server Support for UNIX servers and Midrange servers

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:

   ```
   lsattr -l hdiskX
   ```

   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:
   - Consider upgrading your vendor software and then repeating this procedure. The latest versions of some vendor software might include support for identifying devices using a UDID. For information, see the documentation provided by your vendor software.
   - Put a PVID on the physical volume by running the following command:

     ```
     chdev -l hdiskX -a pv=yes
     ```

Related reference

chdev Command

lspv Command

oem_setup_env Command

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

Find instructions for creating a virtual target device on the Virtual I/O Server.

The following procedure describes how to configure Virtual SCSI. This procedure can be repeated to provide additional virtual disk storage to any client logical partition. This procedure assumes that you already have a physical or logical volume defined on the Virtual I/O Server. For information about physical and logical volumes, see Logical volumes.

This procedure also assumes that the virtual adapters for the Virtual I/O Server and the client partitions were created during the creation of the partition profile. For information about creating the partition, see Installing the Virtual I/O Server.

With the Virtual I/O Server you can export disks as virtual disks. With the Virtual I/O Server you can export two types of physical disks: Virtual SCSI disk backed by a physical volume and a Virtual SCSI disk backed by a logical volume. 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 the virtual target device on the Virtual I/O Server maps the Virtual SCSI adapter with the logical volume or physical disk. Use the **mkvdev** command with the following syntax:

```
mkvdev -vdev TargetDevice -vadapter VirtualSCSIServerAdapter [-dev DeviceName]
```

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
   ent2 Available Virtual I/O Ethernet Adapter (l-lan)
   vhost0 Available Virtual SCSI Server Adapter
   vhost1 Available Virtual SCSI Server Adapter
   vhost2 Available Virtual SCSI Server Adapter
   vsa0 Available LPAR Virtual Serial Adapter
   ```

2. To create a virtual target device, which maps the virtual SCSI server adapter to a physical or logical volume, run the **mkvdev** command. In this procedure, we ran the following command

   ```
   mkvdev –vdev lv_4G –vadapter vhost3
   ```

   In this example, the name of the virtual SCSI server adapter is `vhost3`. The logical volume that was specified was `lv_4G`.

   Note: The **-vdev** flag can specify either a physical or logical volume or an optical device. To map a physical volume to the virtual SCSI server adapter, use `hdiskx` for the **-vdev** flag. For example, if the physical volume name was hdisk5, run `mkvdev -vdev hdisk5 -vadapter vhost3`. To map an optical

device to the virtual SCSI server adapter, use `cdx` for the **-vdev** flag. For example, if the optical device name is cd0, run `mkvdev -vdev cd0 -vadapter vhost3`.
The storage is available to the client 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).

To map a physical volume to the Virtual SCSI Server Adapter, use `hdiskx` instead of the logical volume devices for the -**vdev** flag.

The **lsdev** command shows the newly created Virtual Target Device adapter. For example, running `lsdev -virtual` returns results similar to the following:

```
name status description
vhost0 Available Virtual SCSI Server Adapter
vsa0 Available LPAR Virtual Serial Adapter
vdbsrv Available Virtual Target Device - Logical Volume
```

The **lsmap** command shows the logical connections between newly created devices, as follows:

```
lsmap -vadapter vhost0
```

This command returns results similar to the following:

```
SVSA Physloc Client PartitionID
--------------- -------------------------------------------- ------------------
vhost0 U9111.520.10DDEEC-V1-C20 0x00000000
VTD vdbsrv
LUN 0x8100000000000000
Backing device rootvg_dbsrv
Physloc
```

The physical location is a combination of the slot number, in this case 20, and the logical partition ID. The virtual device can now be attached from the client partition.

Related concepts

"Logical volumes" on page 6
Understand how logical volumes can be exported to client partitions as virtual SCSI disks. A logical volume is a portion of a physical volume.

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

Related tasks

"Installing the Virtual I/O Server manually using the HMC version 6" on page 70
You can create the Virtual I/O Server logical partition and partition profile and install the Virtual I/O Server using the Hardware Management Console (HMC) version 6 or earlier.

Related reference

lsdev Command

lsmap Command

mkvdev Command

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

Find instructions for creating logical volumes and volume groups on the Virtual I/O Server.

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. For conceptual information about logical volumes, see Concepts for virtual SCSI.

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

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

Related concepts

"Concepts for Virtual SCSI" on page 4
Virtual SCSI allows client logical partitions to share disk storage and optical devices that are assigned to the Virtual I/O Server logical partition.

Related reference

mklv Command

mkvg Command

Importing or Exporting a Volume Group

Find instructions for importing and exporting volume groups.

The following procedure explains how to use the import and export procedures to move a user-defined volume group from one system to another. (The rootvg volume group cannot be exported or imported.) The export procedure removes the definition of a volume group from a system. The import procedure introduces the volume group to its new system. You can also use the import procedure to reintroduce a volume group to the system that had been previously associated with and had been exported from. You can also use import and export to add a physical volume that contains data to a volume group by putting the disk to be added in its own volume group.

Note: 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.

To export a volume group, type the following commands:

1. `deactivatevg` *VolumeGroupName*
2. `exportvg` *VolumeGroupName*

To import a volume group, use the importvg command.

Related reference

deactivatevg Command

exportvg Command

importvg Command

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 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 partition ID of the client partition. On the AIX client logical partition, type the following command from the Virtual I/O Server command line:
   a. `oem_setup_env` .
   b. Run `uname -L`

      Your results should look similar to the following:

      ```
      2  fumi02
      ```

      The partition ID is the first number listed. In this example, the partition ID is 2. This number is used in the next step.
   c. Type `exit`.
5. If you have multiple Virtual I/O Server partitions running on your system, determine which Virtual I/O Server 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 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 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.

Related reference

lshwres Command

lsmap Command

lssyscfg Command

oem_setup_env Command

Increasing virtual SCSI device capacity

Increase the size of virtual SCSI disks.

As storage demands increase for virtual client 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 virtual SCSI device capacity, complete the following steps:

1. Increase the size of the physical or logical volumes:
   - Physical volumes: Consult your storage documentation to determine whether your storage subsystem supports expanding the size of a logical unit number (LUN).
   - Logical volumes: Run the extendlv command. For example:
     ```
     extendlv lv3 100M
     ```
     This example increases logical volume *lv3* by 100 MB.

     If there is no additional space in the logical volume, complete the following tasks:
     a. 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*.
     b. Allocate the increased volume to logical partitions by resizing logical volumes. Run the extendlv command to increase the size of a logical volume.
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 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 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 reference

activatevg Command

cfgdev Command

extendlv Command

extendvg Command

Related information

chvg Command

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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. You can change this value from the default value to any value from 1 to 256. The default value is 3. 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 partition, and the maximum transfer size for the virtual SCSI client adapter instance that is the parent device for the disk instance.

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, on the client 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 partition issue the following command:

```
lsattr -E1 hdiskN
```

Related reference

chdev Command

Virtual SCSI reserve and release requirements

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

Complete the following tasks to enable the Virtual I/O Server support of 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. Enable 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.

Related reference

chdev Command

mkvdev Command

Managing users on the Virtual I/O Server

Find commands for creating, listing, changing, switching, and removing users.

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 22. Tasks and associated commands for working with Virtual I/O Server users

Task	Command
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 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

Related reference

chuser Command

lsuser Command

mkuser Command

rmuser Command

su Command

Configuring the IBM Tivoli agents and clients on the Virtual I/O Server

You can configure and start the IBM Tivoli Monitoring agent, IBM Tivoli Usage and Accounting Manager agent, and the IBM Tivoli Storage Manager client.

Related concepts

"IBM Tivoli agents and clients on the Virtual I/O Server" on page 28
Learn about the IBM Tivoli Monitoring agent, IBM Tivoli Storage Manager client, and the IBM Tivoli Usage and Accounting Manager agent packaged with the Virtual I/O Server.

Configuring the IBM Tivoli Monitoring agent

You can configure and start the IBM Tivoli Monitoring agent on the Virtual I/O Server.

IBM Tivoli Monitoring System Edition for System p enables you to monitor the health and availability of multiple IBM System p servers (including the Virtual I/O Server) from the Tivoli Enterprise Portal. IBM Tivoli Monitoring System Edition for System p 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 IBM Tivoli Monitoring.

Before you start, complete the following tasks:

- Ensure that the Virtual I/O Server is running fix pack 8.1.0. For instructions, see Updating the Virtual I/O Server.
- 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 to enable the monitoring agent to 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 or client) 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 concepts

"User types for the Virtual I/O Server" on page 27
Use this topic to provide information about Virtual I/O Server user types and their user permissions.

Related tasks

"Updating the Virtual I/O Server" on page 127
Find instructions for updating the Virtual I/O Server.

"Connecting to the Virtual I/O Server using OpenSSH" on page 57
You can set up remote connections to the Virtual I/O Server using secure connections.

Setting up secure script executions between ssh clients and the HMC

Related reference

Tasks and roles

"Configuration attributes for IBM Tivoli agents and clients" on page 124
Learn about required and optional configuration attributes and variables for the IBM Tivoli Monitoring agent, the IBM Tivoli Usage and Accounting Manager agent, and the IBM Tivoli Storage Manager client.

cfgsvc Command

lssvc Command

startsvc Command

stopsvc Command

Related information

ITM 6.1 documentation

IBM Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide

Configuring the IBM Tivoli Usage and Accounting Manager agent

You can configure and start the IBM 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 IBM Tivoli Usage and Accounting Manager agent on the Virtual I/O Server. IBM 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 used by entities such as cost centers, departments, and users. IBM 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, i5/OS, and VMware.

Before you start, ensure that the Virtual I/O Server is installed. The IBM 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.

To configure and start the IBM Tivoli Usage and Accounting Manager agent, complete the following steps:

1. List all of the available IBM Tivoli Usage and Accounting Manager agents using the **lssvc** command. For example,

   ```
   $lssvc
   ITUAM_base
   ```
2. Based on the output of the **lssvc** command, decide which IBM Tivoli Usage and Accounting Managerg agent you want to configure. For example, `ITUAM_base`
3. List all of the attributes that are associated with the IBM Tivoli Usage and Accounting Manager agent using the **cfgsvc** command. For example:

   ```
   $cfgsvc –ls ITUAM_base
    ACCT_DATA0
   ACCT_DATA1
   ISYSTEM
   IPROCESS
   ```
4. Configure the IBM 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 IBM 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.
5. Start the monitoring agent using the **startsvc** command. For example:

   ```
   startsvc ITUAM_base
   ```

After you start the IBM Tivoli Usage and Accounting Manager agent, it begins to collect data and generate log files. You can configure the IBM Tivoli Usage and Accounting Manager server to retrieve the log files, which are then processed by the IBM Tivoli Usage and Accounting Manager Processing Engine. You can work with the data from the IBM Tivoli Usage and Accounting Manager Processing Engine as follows::

- You can generate customized reports, spreadsheets, and graphs. IBM 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 one of the following resources:

- If you are running the IBM Tivoli Usage and Accounting Manager Processing Engine on Windows, then see the *IBM Tivoli Usage and Accounting Manager Data Collectors for Microsoft Windows User's Guide*.
- If you are running the IBM Tivoli Usage and Accounting Manager Processing Engine on UNIX or Linux, then see the *IBM Tivoli Usage and Accounting Manager Data Collectors for UNIX and Linux User's Guide*.

Related concepts

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

"IBM Tivoli agents and clients on the Virtual I/O Server" on page 28
Learn about the IBM Tivoli Monitoring agent, IBM Tivoli Storage Manager client, and the IBM Tivoli Usage and Accounting Manager agent packaged with the Virtual I/O Server.

Related reference

"Configuration attributes for IBM Tivoli agents and clients" on page 124
Learn about required and optional configuration attributes and variables for the IBM Tivoli Monitoring agent, the IBM Tivoli Usage and Accounting Manager agent, and the IBM Tivoli Storage Manager client.

cfgsvc Command

lssvc Command

startsvc Command

stopsvc Command

Related information

IBM Tivoli Usage and Accounting Manager Data Collectors for UNIX and Linux User's Guide

IBM Usage and Accounting Manager Data Collectors for Microsoft Windows User's Guide

Configuring the IBM Tivoli Storage Manager client

You can configure the IBM Tivoli 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, including IBM System p servers. Configuring the Tivoli Storage Manager client on the Virtual I/O Server enables you to 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.

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. To determine what information you must provide to the Tivoli Storage Manager administrator, see the *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*.

After you are finished, you are ready to back up and restore the Virtual I/O Server using the Tivoli Storage Manager.

Related concepts

"Installing the Virtual I/O Server" on page 50
Find instructions for installing the Virtual I/O Server by deploying a system plan or manually creating the partition and partition profile and installing the Virtual I/O Server.

"IBM Tivoli agents and clients on the Virtual I/O Server" on page 28
Learn about the IBM Tivoli Monitoring agent, IBM Tivoli Storage Manager client, and the IBM Tivoli Usage and Accounting Manager agent packaged with the Virtual I/O Server.

Related tasks

"Backing up the Virtual I/O Server using IBM Tivoli Storage Manager" on page 136
You can use the IBM Tivoli Storage Manager to automatically back up the Virtual I/O Server on regular intervals, or you can perform incremental backups.

"Restoring the Virtual I/O Server using IBM Tivoli Storage Manager" on page 144
You can use the IBM Tivoli Storage Manager to restore the mksysb image of the Virtual I/O Server.

Related reference

"Configuration attributes for IBM Tivoli agents and clients"
Learn about required and optional configuration attributes and variables for the IBM Tivoli Monitoring agent, the IBM Tivoli Usage and Accounting Manager agent, and the IBM Tivoli Storage Manager client.

cfgsvc Command

lssvc Command

Related information

Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide

Configuration attributes for IBM Tivoli agents and clients

Learn about required and optional configuration attributes and variables for the IBM Tivoli Monitoring agent, the IBM Tivoli Usage and Accounting Manager agent, and the IBM Tivoli Storage Manager client.

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 IBM Tivoli Storage Manager or IBM Tivoli Usage and Accounting Manager.

IBM Tivoli Monitoring

Table 23. IBM 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.

Table 23. IBM Tivoli Monitoring configuration attributes (continued)

Attribute	Description
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. 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. 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.

IBM Tivoli Storage Manager

Table 24. 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.

IBM Tivoli Usage and Accounting Manager

Table 25. IBM 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
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. • N: Indicates that the Usage and Accounting Manager agent does not collect accounting data.	Y

Table 25. IBM Tivoli Usage and Accounting Manager configuration variables in the A_config.par file (continued)

Variable	Description	Possible values	Default value
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. • N: Indicates that the Usage and Accounting Manager agent does not collect data about the storage file system.	N

Table 26. IBM 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.

Related tasks

"Configuring the IBM Tivoli Monitoring agent" on page 119
You can configure and start the IBM Tivoli Monitoring agent on the Virtual I/O Server.

"Configuring the IBM Tivoli Usage and Accounting Manager agent" on page 121
You can configure and start the IBM Tivoli Usage and Accounting Manager agent on the Virtual I/O Server.

"Configuring the IBM Tivoli Storage Manager client" on page 123
You can configure the IBM Tivoli Storage Manager client on the Virtual I/O Server.

Related information

ITM 6.1 documentation

IBM Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide

IBM Tivoli Usage and Accounting Manager Data Collectors for UNIX and Linux User's Guide

Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide

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.

Related tasks

"Managing users on the Virtual I/O Server" on page 118
Find commands for creating, listing, changing, switching, and removing users.

Related reference

chuser Command

ldapadd Command

ldapsearch Command

mkldap Command

stopnetsvc Command

Maintaining the Virtual I/O Server

Find information about updating, backing up, restoring, and monitoring the Virtual I/O Server.

Subscribing to the Virtual I/O Server subscription service

A subscription service is available to allow Virtual I/O Server users to stay current on news and product updates.

To subscribe to this service, follow these steps:

1. Go to the Subscription service for UNIX and Linux servers Web site.
2. Click the **Subscribe / Setup** tab and complete the form.

After subscribing, you are notified of all Virtual I/O Server news and product updates.

Related information

Subscription service for UNIX and Linux servers

Updating the Virtual I/O Server

Find instructions for updating the Virtual I/O Server.

In this procedure, you install an update to the Virtual I/O Server. Obtain the update either from a CD that contains the update or download the update. To update the Virtual I/O Server, follow these steps:

1. Make a backup of the Virtual I/O Server by following the steps in Backing up the Virtual I/O Server.
2. Download the required updates from the Virtual I/O Server support site. Alternatively, you can get the updates from the update CD.
3. Install the update using the updateios command. For example, if your update fileset is located in the /home/padmin/update directory, type the following:

   ```
   updateios -dev /home/padmin/update
   ```

Note: The updateios command installs all updates located in the specified directory.

Related tasks

"Backing up the Virtual I/O Server"
You can back up the Virtual I/O Server and user-defined virtual devices to tape, DVD, or a remote file system.You can back up the Virtual I/O Server and user-defined virtual devices using the backupios command. You can also IBM Tivoli Storage Manager to schedule backups and store backups on another server.

Related reference

updateios Command

Related information

Virtual I/O Server Support for UNIX servers and Midrange servers

Backing up the Virtual I/O Server

You can back up the Virtual I/O Server and user-defined virtual devices to tape, DVD, or a remote file system.You can back up the Virtual I/O Server and user-defined virtual devices using the backupios command. You can also IBM 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 27. 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

Related tasks

Backing up partition profile data

"Restoring the Virtual I/O Server" on page 138
You can restore the Virtual I/O Server and user-defined virtual devices from tape, DVD, or a remote file system.You can restore the Virtual I/O Server and user-defined virtual devices using the installios command or IBM Tivoli Storage Manager.

Related reference

backupios Command

installios Command

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. (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.

Related tasks

Backing up and restoring partition data

“Backing up user-defined virtual devices” on page 133
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.

“Restoring the Virtual I/O Server from tape” on page 139
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.

Related reference

backupios Command

bkprofdata Command

cfgdev Command

lsdev Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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. (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 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.

Related tasks

Backing up and restoring partition data

"Backing up user-defined virtual devices" on page 133
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.

"Restoring the Virtual I/O Server from one or more DVDs" on page 140
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.

Related reference

backupios Command

bkprofdata Command

cfgdev Command

lsdev Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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. (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.

Related tasks

Backing up and restoring partition data

"Backing up user-defined virtual devices" on page 133
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.

"Restoring the Virtual I/O Server from the HMC using a nim_resources.tar file" on page 140
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.

Related reference

backupios Command

bkprofdata Command

mkdir Command

mount Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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. (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. To find the latest updates, see the Fix Central Web site.
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.

Related tasks

Backing up and restoring partition data

"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.

"Restoring the Virtual I/O Server from a NIM server using a mksysb file" on page 141
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.

Related reference

backupios Command

bkprofdata Command

mkdir Command

mount Command

Related information

Fix Central

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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
   - Backing up the Virtual I/O Server to one or more DVDs
   - Backing up the Virtual I/O Server to a remote file system by creating a nim_resources.tar file
   - Backing up the Virtual I/O Server to a remote file system by creating a mksysb image
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 28. 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` *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` *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.

Related tasks

"Backing up the Virtual I/O Server to tape" on page 129
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.

"Backing up the Virtual I/O Server to one or more DVDs" on page 130
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.

"Backing up the Virtual I/O Server to a remote file system by creating a nim_resources.tar file" on page 131
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 by creating a mksysb image" on page 132
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.

"Restoring user-defined virtual devices" on page 143
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.

Related reference

activatevg Command

cfgnamesrv Command

entstat Command

hostmap Command

ioslevel Command

lsdev Command

lsuser Command

lsvg Command

netstat Command

optimizenet Command

savevgstruct Command

tee Command

viosecure Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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.

Related tasks

"Backing up user-defined virtual devices" on page 133
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.

Related reference

crontab Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Backing up the Virtual I/O Server using IBM Tivoli Storage Manager:

You can use the IBM Tivoli Storage Manager to automatically back up the Virtual I/O Server on regular intervals, or you can perform incremental backups.

Related tasks

"Configuring the IBM Tivoli Storage Manager client" on page 123
You can configure the IBM Tivoli Storage Manager client on the Virtual I/O Server.

"Restoring the Virtual I/O Server using IBM Tivoli Storage Manager" on page 144
You can use the IBM Tivoli Storage Manager to restore the mksysb image of the Virtual I/O Server.

Backing up the Virtual I/O Server using IBM 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 IBM Tivoli Storage Manager client.
- 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.

   - For instructions about how to save user-defined virtual devices, see Backing up user-defined virtual devices.
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
   ```

Related tasks

"Configuring the IBM Tivoli Storage Manager client" on page 123
You can configure the IBM Tivoli Storage Manager client on the Virtual I/O Server.

"Backing up the Virtual I/O Server to a remote file system by creating a mksysb image" on page 132
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 user-defined virtual devices" on page 133
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.

Related reference

crontab Command

Related information

Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide

Backing up the Virtual I/O Server using IBM 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 IBM Tivoli Storage Manager client.
- 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.
  – For instructions about how to save user-defined virtual devices, see Backing up user-defined virtual devices.

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`.

Related tasks

"Configuring the IBM Tivoli Storage Manager client" on page 123
You can configure the IBM Tivoli Storage Manager client on the Virtual I/O Server.

"Backing up the Virtual I/O Server to a remote file system by creating a mksysb image" on page 132
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 user-defined virtual devices" on page 133
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.

Related information

Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide

Restoring the Virtual I/O Server

You can restore the Virtual I/O Server and user-defined virtual devices from tape, DVD, or a remote file system.You can restore the Virtual I/O Server and user-defined virtual devices using the installios command or IBM 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 29. 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

Table 29. Backup and restoration methods for the Virtual I/O Server (continued)

Backup method	Media	Restoration method
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

Related tasks

"Backing up the Virtual I/O Server" on page 128
You can back up the Virtual I/O Server and user-defined virtual devices to tape, DVD, or a remote file system.You can back up the Virtual I/O Server and user-defined virtual devices using the backupios command. You can also IBM Tivoli Storage Manager to schedule backups and store backups on another server.

Backing up partition profile data

Related reference

installios Command

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. (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 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.

Related tasks

"Backing up the Virtual I/O Server to tape" on page 129
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.

"Restoring user-defined virtual devices" on page 143
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.

Backing up and restoring partition data

Related reference

bootlist Command

rstprofdata Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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. (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.

Related tasks

"Backing up the Virtual I/O Server to one or more DVDs" on page 130
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.

"Restoring user-defined virtual devices" on page 143
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.

Backing up and restoring partition data

Related reference

bootlist Command

rstprofdata Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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. (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.

Related tasks

"Backing up the Virtual I/O Server to a remote file system by creating a nim_resources.tar file" on page 131
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.

"Restoring user-defined virtual devices" on page 143
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.

Backing up and restoring partition data

Related reference

installios Command

rstprofdata Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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. (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. 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 30. 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 tasks

"Backing up the Virtual I/O Server to a remote file system by creating a mksysb image" on page 132
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 and restoring partition data

"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.

Related reference

rstprofdata Command

Related information

nim Command

Using the NIM define operation

Defining a SPOT resource

Installing the Virtual I/O Server using NIM

Installing a client using NIM

IBM System p Advanced POWER Virtualization Best Practices RedPaper

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
- Restoring the Virtual I/O Server from one or more DVDs
- Restoring the Virtual I/O Server from the HMC using a nim_resources.tar file
- Restoring the Virtual I/O Server from a NIM server using a mksysb file

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. Recreate 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`*.

Related tasks

"Backing up user-defined virtual devices" on page 133
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.

"Restoring the Virtual I/O Server from tape" on page 139
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.

"Restoring the Virtual I/O Server from one or more DVDs" on page 140
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.

"Restoring the Virtual I/O Server from the HMC using a nim_resources.tar file" on page 140
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.

"Restoring the Virtual I/O Server from a NIM server using a mksysb file" on page 141
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.

Related reference

lspv Command

mkvdev Command

restorevgstruct Command

tee Command

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Restoring the Virtual I/O Server using IBM Tivoli Storage Manager:

You can use the IBM 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
- Restoring the Virtual I/O Server from one or more DVDs
- Restoring the Virtual I/O Server from the HMC using a nim_resources.tar file
- Restoring the Virtual I/O Server from a NIM server using a mksysb file

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.

Related tasks

"Backing up the Virtual I/O Server using IBM Tivoli Storage Manager" on page 136
You can use the IBM Tivoli Storage Manager to automatically back up the Virtual I/O Server on regular intervals, or you can perform incremental backups.

"Restoring the Virtual I/O Server from tape" on page 139
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.

"Restoring the Virtual I/O Server from one or more DVDs" on page 140
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.

"Restoring the Virtual I/O Server from the HMC using a nim_resources.tar file" on page 140
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.

"Restoring the Virtual I/O Server from a NIM server using a mksysb file" on page 141
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.

Related reference

mkcd Command

mkdvd Command

ftp Command

startnetsvc Command

Related information

AIX Toolbox for Linux Applications

Monitoring the Virtual I/O Server

You can monitor the Virtual I/O Server using error logs or IBM 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 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.

IBM Tivoli Monitoring

With Virtual I/O Server V1.3.0.1 (fix pack 8.1), you can install and configure the IBM Tivoli Monitoring System Edition for System p agent on the Virtual I/O Server. IBM Tivoli Monitoring System Edition for System p enables you to monitor the health and availability of multiple IBM System p servers (including the Virtual I/O Server) from the Tivoli Enterprise Portal. IBM Tivoli Monitoring System Edition for System p 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 IBM Tivoli Monitoring.

Related tasks

"Configuring the IBM Tivoli Monitoring agent" on page 119
You can configure and start the IBM Tivoli Monitoring agent on 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.

This section includes information about troubleshooting the Virtual I/O Server. For information about troubleshooting the Integrated Virtualization Manager, see Troubleshooting with the Integrated Virtualization Manager.

Related tasks

Troubleshooting with the Integrated Virtualization Manager

Troubleshooting the Virtual I/O Server logical partition

Find information and procedures for troubleshooting and diagnosing the Virtual I/O Server 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.

Related tasks

AIX fast-path problem-isolation

Related reference

diagmenu Command

oem_setup_env Command

Viewing statistics for Ethernet drivers and devices

You can use statistical information about virtual Ethernet drivers and devices to troubleshoot network problems.

Display the Ethernet device generic statistics and the Ethernet device-specific statistics for an Ethernet driver or device by running the entstat command. For example:

```
entstat -all ent8
```

The **entstat** command output for a Shared Ethernet Adapter displays the following information:

- Shared Ethernet Adapter statistics:
  - Number of real and virtual adapters (If you are using Shared Ethernet Adapter failover, this number does not include the control channel adapter)
  - Shared Ethernet Adapter flags
  - VLAN IDs
  - Information about real and virtual adapters

- Shared Ethernet Adapter failover statistics:
  - High availability statistics
  - Packet types
  - State of the Shared Ethernet Adapter
  - Bridging mode
- GARP VLAN Registration Protocol (GVRP) statistics:
  - Bridge Protocol Data Unit (BPDU) statistics
  - Generic Attribute Registration Protocol (GARP) statistics
  - GARP VLAN Registration Protocol (GVRP) statistics
- Listing of the individual adapter statistics for the adapters associated with the Shared Ethernet Adapter

Related reference

"Shared Ethernet Adapter statistics"
Learn about general Shared Ethernet Adapter statistics, such as VLAN IDs and packet information, and view examples.

"Shared Ethernet Adapter failover statistics" on page 155
Learn about Shared Ethernet Adapter failover statistics, such as high availability information and packet types, and view examples.

"GARP VLAN Registration Protocol statistics" on page 163
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.

entstat Command

Shared Ethernet Adapter statistics:

Learn about general Shared Ethernet Adapter statistics, such as VLAN IDs and packet information, and view examples.

Statistic descriptions

Table 31. Descriptions of Shared Ethernet Adapter statistics

Statistic	Description
Number of adapters	Includes the real adapter and all of the virtual adapters. **Note:** If you are using Shared Ethernet Adapter failover, then the control channel adapter is not included.

Table 31. Descriptions of Shared Ethernet Adapter statistics (continued)

Statistic	Description
Shared Ethernet Adapter flags	Denotes the features that the Shared Ethernet Adapter is currently running. **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. **LARGESEND** The large send feature has been enabled on the Shared Ethernet Adapter. **JUMBO_FRAMES** The jumbo frames feature has been enabled on the Shared Ethernet Adapter. **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 31. Descriptions of Shared Ethernet Adapter statistics (continued)

Statistic	Description
Real adapters	**Packets received** Number of packets received on the physical network. **Packets bridged** Number of packets received on the physical network that were sent to the virtual network. **Packets consumed** Number of packets received on the physical network that were addressed to the interface configured over the Shared Ethernet Adapter. **Packets fragmented** 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). **Packets transmitted** 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). **Packets dropped** Number of packets received on the physical network that were dropped for one of the following reasons: • The packet was the oldest packet on a thread's queue and there was no space to accommodate a newly received packet. • The packet had an invalid VLAN ID and could not be processed. • The packet was addressed to the Shared Ethernet Adapter interface, but its interface had no filters registered.

Table 31. Descriptions of Shared Ethernet Adapter statistics (continued)

Statistic	Description
Virtual adapters	**Packets received** Number of packets received on the virtual network. In other words, the number of packets received on all of the virtual adapters. **Packets bridged** Number of packets received on the virtual network that were sent to the physical network. **Packets consumed** Number of packets received on the virtual network that were addressed to the interface configured over the Shared Ethernet Adapter. **Packets fragmented** 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. **Packets transmitted** 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). **Packets dropped** Number of packets received on the virtual network that were dropped for one of the following reasons: • The packet was the oldest packet on a thread's queue and there was no space to accommodate a newly received packet. • 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 31. 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
```

Related concepts

"Shared Ethernet Adapters" on page 13
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.

Related reference

entstat Command

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 32. Descriptions of Shared Ethernet Adapter failover statistics

Statistic	Description
High availability	**Control Channel PVID** Port VLAN ID of the virtual Ethernet adapter used as the control channel. **Control Packets in** Number of packets received on the control channel. **Control Packets out** Number of packets sent on the control channel.

Table 32. Descriptions of Shared Ethernet Adapter failover statistics (continued)

Statistic	Description
Packet types	**Keep-Alive Packets** 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. **Recovery Packets** 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. **Notify Packets** 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. **Limbo Packets** 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 32. Descriptions of Shared Ethernet Adapter failover statistics (continued)

Statistic	Description
State	The current state of the Shared Ethernet Adapter. **INIT** The Shared Ethernet Adapter failover protocol has just been initiated. **PRIMARY** The Shared Ethernet Adapter is actively connecting traffic between the VLANs to the network. **BACKUP** The Shared Ethernet Adapter is idle and not connecting traffic between the VLANs and the network. **RECOVERY** The primary Shared Ethernet Adapter recovered from a failure and is ready to be active again. **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. **LIMBO** One of the following situations is true: • The physical network is not operational. • The physical network's state is unknown. • 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. **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. **All** The Shared Ethernet Adapter is sending and receiving all types of network traffic. **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 32. 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. **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. **Standby** 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. **Priority** 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
```

Related concepts

"Shared Ethernet Adapter failover" on page 48
Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption.

Related reference

entstat Command

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 33. Descriptions of BPDU statistics

BPDU statistic	Description
Transmit	**Packets** Number of packets sent. **Failed packets** 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).

Table 33. Descriptions of BPDU statistics (continued)

BPDU statistic	Description
Receive	**Packets** Number of packets received. **Unprocessed Packets** Packets that could not be processed because the protocol was not running at the time. **Non-contiguous Packets** Packets that were received in several packet fragments. **Packets with unknown PID** 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. **Packets with Wrong Length** 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 34. Descriptions of GARP statistics

GARP statistic	Description
Transmit	**Packets** Number of packets sent. **Failed packets** 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). **Leave All Events** Packets sent with event type *Leave All*. **Join Empty Events** Packets sent with event type *Join Empty* **Join In Events** Packets sent with event type *Join In* **Leave Empty Events** Packets sent with event type *Leave Empty* **Leave In Events** Packets sent with event type *Leave In* **Empty Events** Packets sent with event type *Empty*

Table 34. Descriptions of GARP statistics (continued)

GARP statistic	Description
Receive	**Packets** Number of packets received **Unprocessed Packets** Packets that could not be processed because the protocol was not running at the time. **Packets with Unknown Attr Type:** 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). **Leave All Events** Packets received with event type *Leave All* **Join Empty Events** Packets received with event type *Join Empty* **Join In Events** Packets received with event type *Join In* **Leave Empty Events** Packets received with event type *Leave Empty* **Leave In Events** Packets received with event type *Leave In* **Empty Events** 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 35. Descriptions of GVRP statistics

GVRP statistic	Description
Transmit	**Packets** Number of packets sent **Failed packets** 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). **Leave All Events** Packets sent with event type *Leave All*. **Join Empty Events** Packets sent with event type *Join Empty* **Join In Events** Packets sent with event type *Join In* **Leave Empty Events** Packets sent with event type *Leave Empty* **Leave In Events** Packets sent with event type *Leave In* **Empty Events** Packets sent with event type *Empty*

Table 35. Descriptions of GVRP statistics (continued)

GVRP statistic	Description
Receive	**Packets** Number of packets received. **Unprocessed Packets** Packets that could not be processed because the protocol was not running at the time. **Packets with Invalid Length** Packets that contains one or more attributes whose length does not correspond to its event type. **Packets with Invalid Event** Packets that contain one or more attributes whose event type is invalid. **Packets with Invalid Value** Packets that contain one or more attributes whose value is invalid (for example, an invalid VLAN ID). **Total Invalid Attributes** Sum of all of the attributes that had an invalid parameter. **Total Valid Attributes** Sum of all of the attributes that had no invalid parameters. **Leave All Events** Packets sent with event type *Leave All*. **Join Empty Events** Packets sent with event type *Join Empty* **Join In Events** Packets sent with event type *Join In* **Leave Empty Events** Packets sent with event type *Leave Empty* **Leave In Events** Packets sent with event type *Leave In* **Empty Events** 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
```

Related concepts

"Shared Ethernet Adapters" on page 13
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.

Related reference

entstat Command

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 configuraiton 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.

Related reference

chdev Command

entstat Command

lsdev Command

netstat Command

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 partition can ping another client partition on the same system without going through the Virtual I/O Server. If this fails, the problem is likely in the client partition's virtual Ethernet setup. If the ping is successful, proceed to the next step.
2. Start a ping on the source 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 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. See Customer service, support, and troubleshooting for information about service and support.
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.

Related concepts

Customer service, support, and troubleshooting

Related tasks

"Viewing statistics for Ethernet drivers and devices" on page 147
You can use statistical information about virtual Ethernet drivers and devices to troubleshoot network problems.

Related reference

entstat Command

starttrace Command

stoptrace Command

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.

Related reference

startnetsvc Command

Troubleshooting the client logical partition

Find information and procedures for troubleshooting the client 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 36. Labels, identifiers, error numbers, problem descriptions, and resolutions of common Virtual SCSI client partition problems

Label	Identifier	Error Number	Problem	Resolution
VSCSI_ERR2	857033C6	0006 RC 0000 0002	The Virtual SCSI server adapter on the Virtual I/O Server partition is not open.	Make the server adapter on the Virtual I/O Server partition available for use.
		001C RC 0000 0000	The Virtual SCSI server adapter on the Virtual I/O Server partition has been closed abruptly.	Determine why the server adapter in the Virtual I/O Server partition was closed.
VSCSI_ERR3	ED995F18	000D RC FFFF FFF0	The Virtual SCSI server adapter on the Virtual I/O Server partition is being used by another client.	Terminate the client partition that is using the server adapter.
		000D RC 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.

Recovering from disks not displaying in SMS

Learn how to recover from disks not displaying in the System Management Services (SMS) menu 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 Obtaining HMC machine code updates and upgrades .
   b. To check whether you have the latest firmware, see Obtaining firmware updates .
   c. To check whether you have the latest level of the Virtual I/O Server, see Updating the Virtual I/O Server.
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:
   - Instructions for HMC version 7 or later:
     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.

- Instructions for HMC version 6 or earlier:
  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. 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.
5. The **Adapter** value on the client logical partition must match the **Remote Adapter** on the Virtual I/O Server partition, and the **Adapter** value on the Virtual I/O Server 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.
6. Verify that the mappings are correct based on your HMC version:
   - HMC version 7 or later: The **Adapter** value on the client logical partition must match the **Remote Adapter** on the Virtual I/O Server partition, and the **Adapter** value on the Virtual I/O Server 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.
   - HMC version 6 or earlier: The **Slot number** value on the client logical partition must match the **Remote partition virtual slot number** on the Virtual I/O Server partition, and the **Slot number** value on the Virtual I/O Server partition must match the **Remote partition virtual slot number** on the client logical partition. If these numbers do not match, from the HMC, modify the profile properties to reflect the correct mapping.
7. The **Slot number** value on the client logical partition must match the **Remote partition virtual slot number** on the Virtual I/O Server partition, and the **Slot number** value on the Virtual I/O Server partition must match the **Remote partition virtual slot number** on the client logical partition. If these numbers do not match, from the HMC, modify the profile properties to reflect the correct mapping.
8. From the Virtual I/O Server command line, type `cfgdev`.
9. Shut down and reactivate the client logical partition.
10. 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 partition ID is 2 (`0x00000002`).

Note: If the client 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.

11. 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
```

Related tasks

Getting HMC machine code fixes

Getting server firmware and power subsystem firmware fixes

"Updating the Virtual I/O Server" on page 127
Find instructions for updating the Virtual I/O Server.

Related reference

cfgdev Command

lsdev Command

lsmap Command

Related information for Using the Virtual I/O Server

IBM Redbooks® (in PDF format), Web sites, and information center topics contain information related to the Using the Virtual I/O Server topic. You can view or print any of the PDF files.

Web sites

- Advanced POWER Virtualization Web site (http://www-03.ibm.com/systems/p/apv/index.html). The Advanced POWER Virtualization Web site provides information and resources for the Virtual I/O Server, Integrated Virtualization Manager, Partition Load Manager, and Micro-Partitioning.
- IBM System p and AIX Information Center Web site (http://publib.boulder.ibm.com/infocenter/pseries) The IBM System p and AIX Information Center is a source for technical information about IBM System p servers and AIX. The Information Center is your starting point for all System p technical information.
- IBM System Planning Tool Web site (http://www.ibm.com/systems/support/tools/systemplanningtool/) You can download the System Planning Tool from the IBM System Planning Tool Web site.
- Linux servers Web site (http://www.ibm.com/eserver/linux) The Linux server Web site is a source for technical information about Linux on IBM Systems servers.
- Tivoli software information center Web site (http://publib.boulder.ibm.com/tividd/td/link/tdprodlist.html) The Tivoli software information center Web site provides an alphabetical list of current Tivoli products. Each product listed provides a link to its documentation.
- Virtual I/O Server Web site (http://techsupport.services.ibm.com/server/virtualization/vios) The Virtual I/O Server Web site provides product updates and additional information.

Redbooks

- Advanced POWER Virtualization on IBM System p5
- IBM System p Advanced POWER Virtualization Best Practices

- LPAR Simplification Tools Handbook
- Partitioning Implementations for IBM eServer p5 Servers

Other information

- Customer service and support
- Partitioning with the Integrated Virtualization Manager
- Managing your server

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

IBM and accessibility

See the IBM Accessibility Center at http://www.ibm.com/able/ for more information about the commitment that IBM has to accessibility.

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
eServer
HACMP
IBM
POWER
PowerVM
pSeries
Redbooks
System p
System p5
Tivoli
Tivoli Enterprise

Microsoft, Windows, Windows NT[®], and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a 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.

Printed in USA