

Hardware Information

Partitioning the server

ESCALA POWER5



REFERENCE
86 A1 18EW 00

ESCALA POWER5

Hardware Information

Partitioning the server

Hardware

July 2006

BULL CEDOC

357 AVENUE PATTON

B.P.20845

49008 ANGERS CEDEX 01

FRANCE

REFERENCE

86 A1 18EW 00

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

Copyright © Bull SAS 1992, 2006

Printed in France

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

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

Trademarks and Acknowledgements

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

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

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

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

Table of Contents

Partitioning the server	1
What's new for Partitioning the server.....	1
October 2005.....	1
Printable PDF.....	2
Saving PDF files.....	2
Concepts for partitioning the server.....	3
Logical partition overview.....	3
Partitioning tools.....	5
Logical partition functional differences.....	11
Types of logical partition configurations.....	17
Minimum hardware configuration requirements for logical partitions.....	18
Physical and virtual hardware resources.....	19
Communications options for logical partitions.....	29
Logical partition utilities.....	31
Partitioning with an HMC.....	32
Creating partitions from a system plan.....	33
Related information for logical partitions.....	38
Related information for AIX logical partitions.....	39
Related information for Linux logical partitions.....	39
Related information for the Integrated Virtualization Manager.....	40
Related information for the Virtual I/O Server.....	40

Partitioning the server

Partitioning the server

Logical partitions enable you to distribute resources within a single server to make it function as if it were two or more independent servers. Plan your next upgrade to include logical partitions.

This information explains how to set up, manage, and troubleshoot AIX and Linux logical partitions.

- **[What's new for Partitioning the server](#)**
This topic highlights changes made with this release that relate to partitioning in general.
 - **[Printable PDF](#)**
Use this to view and print a PDF of this information.
 - **[Concepts for partitioning the server](#)**
Before you start creating logical partitions, it is essential that you understand the concepts behind this type of system configuration.
 - **[Partitioning with an HMC](#)**
The *Hardware Management Console (HMC)* is a system that controls managed systems, including the management of logical partitions and use of Power On Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information for analysis.
 - **[Partitioning with the Integrated Virtualization Manager](#)**
The Integrated Virtualization Manager is a browser-based system management interface that you can use to manage a single managed system that uses Virtual I/O Server on a management partition. You can use the Integrated Virtualization Manager to create logical partitions on a single managed system, manage the virtual storage and virtual Ethernet on the managed system, and view service information related to the managed system.
 - **[Partitioning using the Virtual I/O Server](#)**
The Virtual I/O Server provides virtual storage and shared Ethernet capability to client logical partitions on the system. It allows physical adapters with attached disks on the Virtual I/O Server logical partition to be shared by one or more client partitions. Virtual I/O Server partitions are not intended to run applications or for general user logins. The Virtual I/O Server is installed in its own logical partition.
 - **[Related information for logical partitions](#)**
There are Web sites, and ESCALA Power5 Hardware Information topics that relate to the Partitioning the server topic.
-

What's new for Partitioning the server

This topic highlights changes made with this release that relate to partitioning in general.

October 2005

Integrated Virtualization Manager

The Integrated Virtualization Manager is a browser-based system management interface for Virtual I/O Server that allows you to create and manage AIX and Linux logical partitions on a single server. See [Partitioning with the Integrated Virtualization Manager](#) for more information.

Virtual Partition Manager

The Virtual Partition Manager is a feature of V5R3 that allows you to create and manage up to four Linux logical partitions from a single logical partition on an model server. See [Partitioning with the Virtual Partition Manager](#) for more information.

InfiniBand (IB)

You can now use InfiniBand (IB) devices with logical partitions on systems that are managed using a Hardware Management Console (HMC). You can assign a globally unique ID (GUID) on an InfiniBand (IB) adapter to a partition profile as required, and the logical partition can use the device associated with that GUID when you activate the logical partition using the partition profile. You can assign only one GUID from each physical InfiniBand (IB) adapter to each partition profile. Also, each GUID on an InfiniBand (IB) adapter can be used by only one logical partition at a time.

Operating system enhancements

- [What's new for Partitioning for AIX](#)
- [What's new for Partitioning for Linux](#)

Parent topic: [Partitioning the server](#)

Printable PDF

Use this to view and print a PDF of this information.

Partitioning with an HMC:

- [Partitioning for AIX with an HMC](#) (about 470 KB) contains the following topics:
 - ◆ What's new for Partitioning for AIX with an HMC
 - ◆ General concepts for partitioning the server
 - ◆ Scenarios for AIX logical partitions
 - ◆ Creating partitions from a system plan
 - ◆ Configuring AIX logical partitions
 - ◆ Configuring resource management for AIX partitions with the Partition Load Manager
 - ◆ Managing AIX logical partitions
 - ◆ Using AIX installed on a logical partition
 - ◆ Troubleshooting AIX logical partitions
 - ◆ Related information for AIX logical partitions
- [Partitioning for Linux with an HMC](#) (about 926 KB) contains the following topics:
 - ◆ What's new for Partitioning for Linux with an HMC
 - ◆ General concepts for partitioning the server
 - ◆ Scenarios for Linux logical partitions
 - ◆ Supported virtual devices for Linux logical partitions
 - ◆ Configuring Linux logical partitions
 - ◆ Managing Linux logical partitions
 - ◆ Using Linux installed on a logical partition
 - ◆ Troubleshooting Linux logical partitions
 - ◆ Related information for Linux logical partitions
- [Partitioning using the Integrated Virtualization Manager](#) (about 320 KB) contains the following topics:
 - ◆ Setting up the initial logical partition
 - ◆ Viewing and modifying system properties
 - ◆ Modifying virtual Ethernet settings
 - ◆ Managing Partitions
 - ◆ Managing storage devices
 - ◆ Creating and modifying user accounts
 - ◆ Troubleshooting with the Integrated Virtualization Manager
 - ◆ Related information for the Integrated Virtualization Manager
- [Partitioning using the Virtual I/O Server](#) (about 1024 KB) contains the following topics:
 - ◆ Virtual I/O Server overview
 - ◆ Concepts for the Virtual I/O Server
 - ◆ Planning for the Virtual I/O Server
 - ◆ Installing and configuring the Virtual I/O Server
 - ◆ Scenarios for the Virtual I/O Server
 - ◆ Managing the Virtual I/O Server
 - ◆ Troubleshooting the Virtual I/O Server
 - ◆ Related information for the Virtual I/O Server

Saving PDF files

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

1. Right-click the PDF link in your browser.

2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you want to save the PDF.
4. Click Save.

Downloading Adobe Reader

You need Adobe Reader installed on your system to view or print these PDFs. You can download a free copy from the [Adobe Web site](http://www.adobe.com/products/acrobat/readstep2.html) (www.adobe.com/products/acrobat/readstep2.html).

Parent topic: [Partitioning the server](#)

Concepts for partitioning the server

Before you start creating logical partitions, it is essential that you understand the concepts behind this type of system configuration.

The purpose of this information is to familiarize you with the hardware and software required for logical partitions and to prepare you to plan for and create logical partitions on your server.

- **Logical partition overview**
This topic summarizes the purpose, process, and tools available to create and manage logical partitions on and ESCALA PL Series hardwareserver hardware.
- **Partitioning tools**
You must use tools to partition your servers. The tool that you use to partition each server depends upon the server model and the operating systems and features that you want to use on the server.
- **Logical partition functional differences**
New server hardware provides several technical enhancements in logical partition functionality compared to previous hardware models and between different partition management tools.
- **Types of logical partition configurations**
This topic provides an overview of the types of logical partition configurations including the manufacturing default configuration, and the Virtual I/O Server partition.
- **Minimum hardware configuration requirements for logical partitions**
This topic provides an overview of minimum hardware configuration requirements for AIX and Linux logical partitions.
- **Physical and virtual hardware resources**
This topic provides an overview of the physical and virtual hardware resources used for AIX and Linux logical partitions including virtual adapters, processors, memory, expansion units, console options, I/O devices, and how each operating system implements virtual resources.
- **Communications options for logical partitions**
The types of communications options that you use depend on your business needs and the operating system you are running. You can establish interpartition communication using virtual Ethernet.
- **Logical partition utilities**
This topic provides an overview of the utilities available to manage logical partitions including the Partition Load Manager for AIX.

Parent topic: [Partitioning the server](#)

Logical partition overview

This topic summarizes the purpose, process, and tools available to create and manage logical partitions on and ESCALA PL Series hardwareserver hardware.

Logical partitioning is the ability to make a server run as if it were two or more independent servers. When you logically partition a server, you divide the resources on the server into subsets called *logical partitions*. Processors, memory, and input/output devices are examples of resources that you can assign to logical partitions. You can install software on the logical partition, and the logical partition runs as an independent

logical server with the processor, memory, and I/O resources that you have allocated to it. Examples of the kinds of software that you can install and run on logical partitions include the AIX and Linux operating systems and Virtual I/O Server software.

Logical partitions share a few system attributes, such as the system serial number, system model, and processor feature code. All other system attributes can vary from one logical partition to another.

You must use tools to partition your servers. The tool that you use to partition each server depends upon the server model and the operating systems and features that you want to use on the server.

Benefits of partitioning

The following scenarios illustrate the benefits of partitioning your server:

Consolidating servers

A logically partitioned server can reduce the number of servers that are needed within an enterprise. You can consolidate several servers into a single logically partitioned system. This eliminates the need for, and expense of, additional equipment.

Sharing resources

You can quickly and easily move hardware resources from one logical partition to another as needs change. Features such as Micro-Partitioning allow for processor resources to be shared automatically among logical partitions that use the shared processor pool. Other features, such as dynamic logical partitioning, allow for resources to be moved to, from, and between running logical partitions manually without shutting down or restarting the logical partitions.

Maintaining independent servers

Dedicating a portion of the resources (disk storage unit, processors, memory, and I/O devices) to a partition achieves logical isolation of software. If configured properly, logical partitions also have some hardware fault tolerance.

Creating a mixed production and test environment

You can create a combined production and test environment on the same server. The production partition can run your main business applications, and the test partition is used to test software. A failure in a test partition, while not necessarily planned, will not disrupt normal business operations.

Merging production and test environments

Partitioning enables separate partitions to be allocated for production and test servers, eliminating the need to purchase additional hardware and software. When testing has been completed, the resources allocated to the test partition can be returned to the production partition or elsewhere as required. As new projects are developed, they can be built and tested on the same hardware on which they will eventually be deployed.

Running integrated clusters

Using high-availability application software, your partitioned server can run as an integrated cluster. You can use an integrated cluster to protect your server from most unscheduled failures within a partition.

Although there are many benefits to using logical partitioning, consider the following points before choosing to partition.

- Processor and memory failures might result in the failure of the entire server with all of its logical partitions. (The failure of a single I/O device affects only the logical partition to which the I/O device belongs.) To reduce the possibility of system failure, you can use the Advanced System Management Interface (ASMI) to set the server so that the server can deconfigure failing processors or memory modules automatically. After the server deconfigures the failing processor or memory module, the server continues running without using the deconfigured processor or memory module.
- There are many concepts that you must understand to implement logical partitions successfully on your server.
- Administering a consolidated system might be more difficult in some ways than administering multiple smaller systems, particularly if the resources in the consolidated system are used at a level close to their capacity. If you anticipate that you will use your server at a level close to its capacity, consider ordering a server model that is capable of Power On Demand (POD).

Sharing resources

Although each logical partition acts as an independent server, the logical partitions on a server can share some kinds of resources with each other. The ability to share resources among many logical partitions allows you to increase resource utilization on the server and to move the server resources to where they are needed. The following list illustrates some of the ways in which logical partitions can share resources. (For some server models, the features mentioned in this list are options for which you must obtain and enter an activation code.)

- Micro-Partitioning (or shared processing) allows logical partitions to share the processors in the shared processor pool. The shared processor pool includes all processors on the server that are not dedicated to specific logical partitions. Each logical partition that uses the shared processor pool is assigned a specific amount of processor power from the shared processor pool. If the logical partition needs more processor power than its assigned amount, the logical partition is set by default to use the unused processor power in the shared processor pool. The amount of processor power that the logical partition can use is limited only by the virtual processor settings of the logical partition and the amount of unused processor power available in the shared processor pool. For more information about Micro-Partitioning, see [Shared processors](#).
- Dynamic logical partitioning allows you to move resources to, from, and between running logical partitions manually without shutting down or restarting the logical partitions. This allows you to share devices that logical partitions use occasionally. For example, if the logical partitions on your server use an optical drive occasionally, you can assign a single optical drive to multiple logical partitions as a desired device. The optical drive would belong to only one logical partition at a time, but you can use dynamic logical partitioning to move the optical drive between logical partitions as needed. On servers that are managed using the Integrated Virtualization Manager, dynamic logical partitioning is supported only for the management partition. Dynamic logical partitioning is not supported on servers that are managed using the Virtual Partition Manager.
- Virtual I/O allows logical partitions to access and use I/O resources on other logical partitions. For example, virtual Ethernet allows you to create a virtual LAN that connects the logical partitions on your server to each other. If one of the logical partitions on the server has a physical Ethernet adapter that is connected to an external network, you can configure the operating system of that logical partition to connect the virtual LAN with the physical Ethernet adapter. This allows the logical partitions on the server to share a physical Ethernet connection to an external network.

Supported operating systems and software

The AIX and Linux operating systems can be installed on specific server models. Also, Virtual I/O Server software can be installed on specific server models. Virtual I/O Server provides virtual I/O resources to other logical partitions with general-purpose operating systems.

Managed Systems

You can install a single operating system on a managed system and use the managed system as a single server. Alternately, you can use a partitioning tool, such as the Hardware Management Console (HMC), to create multiple logical partitions on the managed system. The partitioning tool manages the logical partitions on the managed system.

Parent topic: [Concepts for partitioning the server](#)

Partitioning tools

You must use tools to partition your servers. The tool that you use to partition each server depends upon the server model and the operating systems and features that you want to use on the server.

You can use the following tools to partition your server:

- Hardware Management Console (HMC)
- Integrated Virtualization Manager

- **Hardware Management Console**

The *Hardware Management Console (HMC)* is a system that controls managed systems, including the management of logical partitions and use of Power On Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information for analysis.

- **Integrated Virtualization Manager**

The *Integrated Virtualization Manager* is a browser-based system management interface for the Virtual I/O Server that allows you to create and manage AIX and Linux logical partitions on a single server.

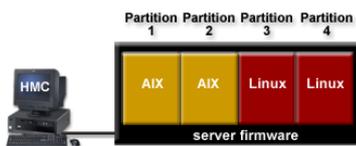
- **Virtual Partition Manager**

The *Virtual Partition Manager* is a feature of that allows you to create and manage one logical partition and up to four Linux logical partitions on a single model server.

Parent topic: [Concepts for partitioning the server](#)

Hardware Management Console

The *Hardware Management Console (HMC)* is a system that controls managed systems, including the management of logical partitions and use of Power On Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information for analysis.



In this figure, you can see the logical partitions and the server firmware on the server. The *server firmware* is code that is stored in system flash memory on the server. The server firmware directly controls the resource allocations on the server and the communications between logical partitions on the server.

If you use a single HMC to manage a server, and the HMC malfunctions or becomes disconnected from the server firmware, then the server continues to run, but you will not be able to change the logical partition configuration of the server. If desired, you can attach an additional HMC to act as a backup and to provide a redundant path between the server and service and support.

Partitioning using the HMC is supported on all server models. Some models require you to enter an activation code before you can partition the server.

- **Partition profile**

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a partition profile, the managed system attempts to start the logical partition using the configuration information in the partition profile. This topic includes a flash demo that shows an example scenario of creating partition profiles on server hardware.

- **Partition profiles that use all of the system resources**

You can create partition profiles on your Hardware Management Console (HMC) that specify all of the resources on the managed system. If you activate such a partition profile when no logical partitions are running on the managed system, then the managed system assigns all of its resources to the logical partition to which the partition profile belongs.

- **System profile**

A *system profile* is an ordered list of partition profiles that is used by the Hardware Management Console (HMC) to start the logical partitions on a managed system in a specific configuration. This topic includes a flash demo that shows an example scenario of creating system profiles on server hardware.

Parent topic: [Partitioning tools](#)

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a partition profile, the managed system attempts to start the logical partition using the configuration information in the partition profile. This topic includes a flash demo that shows an example scenario of creating partition profiles on server hardware.

A partition profile specifies the desired system resources for the logical partition and the minimum and maximum amounts of system resources that the logical partition can have. The system resources specified within a partition profile includes processors, memory, and I/O resources. The partition profile can also specify certain operating settings for the logical partition. For example, you can set a partition profile so that, when the partition profile is activated, the logical partition is set to start automatically the next time that you power on the managed system.

Each logical partition on a managed system that is managed by an HMC has at least one partition profile. If desired, you can create additional partition profiles with different resource specifications for your logical partition. If you create multiple partition profiles, you can designate any partition profile on the logical partition to be the default partition profile. The HMC activates the default profile if you do not select a specific partition profile to be activated. Only one partition profile can be active at one time. To activate another partition profile for a logical partition, you must shut down the logical partition before you activate the other partition profile.

A partition profile is identified by partition ID and profile name. Partition IDs are whole numbers used to identify each logical partition that you create on a managed system, and profile names identify the partition profiles that you create for each logical partition. Each partition profile on a logical partition must have a unique profile name, but you can use a profile name for different logical partitions on a single managed system. For example, logical partition 1 cannot have more than one partition profile with a profile name of `normal`, but you can create a `normal` partition profile for each logical partition on the managed system.

When you create a partition profile, the HMC shows you all of the resources available on your system. The HMC does not verify if another partition profile is currently using a portion of these resources. Therefore, it is possible for you to overcommit resources. When you activate a profile, the system attempts to allocate the resources that you assigned to the profile. If you have overcommitted resources, the partition profile will not be activated.

For example, you have four processors on your managed system. Partition 1 profile A has three processors, and partition 2 profile B has two processors. If you attempt to activate both of these partition profiles at the same time, partition 2 profile B will fail to activate because you have overcommitted processor resources.

When you shut down a logical partition and reactivate the logical partition using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition using dynamic logical partitioning are lost when you reactivate the logical partition using a partition profile. This is desirable when you want to undo dynamic logical partitioning changes to the logical partition. However, this is not desirable if you want to reactivate the logical partition using the resource specifications that the logical partition had when you shut down the managed system. It is therefore best to keep your partition profiles up to date with the latest resource specifications. You can save the current configuration of the logical partition as a partition profile. This allows you to avoid having to change partition profiles manually. For more information about this procedure, see [Saving the partition configuration to a partition profile](#).

If you shut down a logical partition whose partition profiles are not up to date, and the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system using the partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.

Memory and processor resource assignment

When you create a partition profile for a logical partition, you set up the desired, minimum, and maximum amounts of memory and processor resources that you want for the logical partition. The desired value is the resource amount that the logical partition gets if you do not overcommit the resource on the managed system. If the desired amount of resources is available when you activate the partition profile, then the logical partition starts with the desired amount of resources. However, if the desired amount of resources is not available when you activate the partition profile, then the resources on your managed system are overcommitted. In that case, if the amount of resources that are available on the managed system is equal to or greater than the

minimum amount of resources in the partition profile, then the logical partition starts with the available amount of resources. If the minimum amount of resources is not met, then the logical partition does not start.

Virtual processor assignment

If you create a partition profile that is set to use shared processors, the HMC calculates a minimum, maximum, and desired number of virtual processors for the partition profile. The calculation is based upon the minimum, maximum, and desired number of processing units that you specify for the partition profile. By default, the virtual processor settings are calculated as follows:

- The default minimum number of virtual processors is the minimum number of processing units (rounded up to the next whole number). For example, if the minimum number of processing units is 0.8, the default minimum number of virtual processors is 1.
- The default desired number of virtual processors is the desired number of processing units (rounded up to the next whole number). For example, if the desired number of processing units is 2.8, the default desired number of virtual processors is 3.
- The default maximum number of virtual processors is the maximum number of processing units rounded up to the next whole number and multiplied by two. For example, if the maximum number of processing units is 3.2, the default maximum number of virtual processors is 8 (4 times 2).

When you activate the logical partition using the partition profile on the HMC, the logical partition is assigned the desired number of virtual processors. You can then use dynamic logical partitioning to change the number of virtual processors to any number between the minimum and maximum values, so long as the number of virtual processors is greater than the number of processing units that are assigned to the logical partition. Before changing the default settings, performance modeling should be performed.

For example, you create a partition profile on the HMC with the following processor unit settings.

- Minimum processing units 1.25
- Desired processing units 3.80
- Maximum processing units 5.00

The default virtual processor settings for this partition profile on the HMC are as follows.

- Minimum virtual processors 2
- Desired virtual processors 4
- Maximum virtual processors 10

When you activate the logical partition using this partition profile on the HMC, the operating system sees four processors, because the logical partition is activated with the desired value of four virtual processors. Each of these virtual processors has 0.95 processing units supporting the work assigned to the processor. After the logical partition is activated, you can use dynamic logical partitioning to change the number of virtual processors on the logical partition to any number between 2 and 10, so long as the number of virtual processors is greater than the number of processing units that are assigned to the logical partition. If you increase the number of virtual processors, bear in mind that you will have less processing power supporting the work assigned to each processor.

I/O device assignment

I/O devices are assigned to partition profiles on a slot-by-slot basis. Most I/O devices can be assigned to a partition profile on the HMC as required or as desired.

- If an I/O device is assigned to a partition profile as required, then the partition profile cannot be successfully activated if the I/O device is unavailable or is in use by another logical partition. Also, after the logical partition starts, you cannot use dynamic logical partitioning to remove the required I/O device from the running logical partition or move the required I/O device to another logical partition. You must shut down the logical partition before the I/O device can be used by another logical partition. This setting is suitable for devices that are required for the continuous operation of the logical partition (such as disk drives).

- If an I/O device is assigned to a partition profile as desired, then the partition profile can be successfully activated if the I/O device is unavailable or is in use by another logical partition. The desired I/O device can also be removed from the running logical partition or moved to another logical partition using dynamic logical partitioning. This setting is suitable for devices that you want to share among multiple logical partitions (such as optical drives or tape drives).

The exception to this rule is InfiniBand (IB) adapters, which are always added to partition profiles on the HMC as required. Each physical InfiniBand (IB) adapter contains a set of 64 globally unique IDs (GUIDs) that can be assigned to partition profiles. You can assign multiple GUIDs to each partition profile, but you can assign only one GUID from each physical InfiniBand (IB) adapter to each partition profile. Also, each GUID can be used by only one logical partition at a time. You can create multiple partition profiles with the same GUID, but only one of those partition profiles can be activated at a time.

Parent topic: [Hardware Management Console](#)

Partition profiles that use all of the system resources

You can create partition profiles on your Hardware Management Console (HMC) that specify all of the resources on the managed system. If you activate such a partition profile when no logical partitions are running on the managed system, then the managed system assigns all of its resources to the logical partition to which the partition profile belongs.

If you add additional resources to the managed system, the managed system automatically assigns the added resources to the logical partition when the profile is activated. The profile must be activated while the server is in 'partition standby' state, because automatic restart of the partition will not assign newly added processor and memory resources. You do not need to change the partition profile for the managed system to assign the additional resources to the logical partition.

Because all of the hardware is assigned to the logical partition as required, no other logical partitions can be started when the partition profile that uses all of the system resources is running. Likewise, the partition profile that uses all of the system resources cannot be activated if any logical partition is running.

Parent topic: [Hardware Management Console](#)

System profile

A *system profile* is an ordered list of partition profiles that is used by the Hardware Management Console (HMC) to start the logical partitions on a managed system in a specific configuration. This topic includes a flash demo that shows an example scenario of creating system profiles on server hardware.

When you activate the system profile, the managed system attempts to activate each partition profile in the system profile in the order specified. A system profile helps you activate or change the managed system from one complete set of logical partition configurations to another.

It is possible for you to create a system profile that has a partition profile that has overcommitted resources. You can use the HMC to validate the system profile against the currently available system resources and against the total system resources. Validating your system profile ensures that your I/O devices and processing resources are not overcommitted, and it increases the likelihood that the system profile can be activated. The validation process estimates the amount of memory needed to activate all of the partition profiles in the system profile. It is possible that a system profile can pass validation and yet not have enough memory to be activated.

Parent topic: [Hardware Management Console](#)

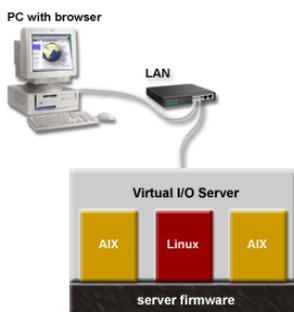
Integrated Virtualization Manager

The *Integrated Virtualization Manager* is a browser-based system management interface for the Virtual I/O Server that allows you to create and manage AIX and Linux logical partitions on a single server.

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 the exact amount of memory and processor resources that the logical partition is to get. 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, but you can change the minimum and maximum values using the Integrated Virtualization Manager if the logical partition is not running.

When you use the Integrated Virtualization Manager to partition your managed system, a fraction of the memory and 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

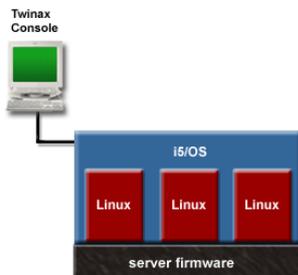
these storage pools and assigned to logical partitions. Physical Ethernet connections are generally shared by configuring the physical Ethernet adapter as a virtual Ethernet bridge between the virtual LAN on the server and an external, physical LAN.

Parent topic: [Partitioning tools](#)

Virtual Partition Manager

The *Virtual Partition Manager* is a feature of that allows you to create and manage one logical partition and up to four Linux logical partitions on a single model server.

To use the Virtual Partition Manager, you must first install on an unpartitioned server. After you install , you can initiate a console session on and use Dedicated Service Tools (DST) or System Service Tools (SST) to create and configure Linux logical partitions. controls the resource allocations of the logical partitions on the server.



This figure illustrates the logical partition and the Linux logical partitions that are managed by the logical partition. The twinaxial console connects to the logical partition, and you can use DST or SST to create and configure the Linux logical partitions on the server.

When you use the Virtual Partition Manager to partition an model server, DST and SST are the only tools that you can use to create and manage the logical partitions. You cannot use Navigator to create or manage logical partitions on an model server. However, the console session that you use to access DST or SST can be initiated using either Operations Console (LAN or direct attach) or a twinaxial console.

Parent topic: [Partitioning tools](#)

Logical partition functional differences

New server hardware provides several technical enhancements in logical partition functionality compared to previous hardware models and between different partition management tools.

- **[AIX logical partition functional differences between new and previous hardware models](#)**
This topic describes logical partition (LPAR) functional differences between new hardware and previous hardware models.
- **[Linux logical partition functional differences between server hardware and previous hardware models](#)**
Use this information to see the new or significantly changed enhancements to Linux partitioning.

Parent topic: [Concepts for partitioning the server](#)

AIX logical partition functional differences between new and previous hardware models

This topic describes logical partition (LPAR) functional differences between new hardware and previous hardware models.

To identify AIX logical partition technical enhancements running on new hardware compared to previous hardware models, see the following table:

Logical partition (LPAR) function	AIX 5L on previous server models	AIX 5L on new server models
LPAR user interface	Hardware Management Console (HMC) required to partition the server.	<ul style="list-style-type: none"> • Hardware Management Console (HMC) • Integrated Virtualization Manager
LPAR authority	User roles created using the HMC: <ul style="list-style-type: none"> • System administrator • Advanced operator • Service representative • Operator • User administrator • Viewer 	Tasks and roles created using the HMC: <ul style="list-style-type: none"> • Super administrator • Operator • Viewer • Product engineer • Service representative User roles created using the Integrated Virtualization Manager: <ul style="list-style-type: none"> • padmin • View/Modify • View Only • Development engineer • Service representative
Maximum number of logical partitions	The maximum number of partitions supported depends on the number of processors in the server model.	The maximum number of partitions supported on systems managed by an HMC is 254 but some server models might support fewer. The maximum number of logical partitions on systems managed by the Integrated Virtualization Manager is 10 times the number of processors.
Types of partitions	Profiles are specific to the HMC. <ul style="list-style-type: none"> • Partition profiles • System profiles 	Profiles are specific to the HMC. <ul style="list-style-type: none"> • Partition profiles • System profiles On systems managed by the Integrated Virtualization Manager, the Virtual I/O Server logical partition is called the management partition.
Processors	Dynamic: can be changed without restarting the partition.	Dynamic: can be changed without restarting the partition.
Memory	<ul style="list-style-type: none"> • Dynamic: can be changed without restarting the partition. • Memory assigned in increments of 256 MB. 	<ul style="list-style-type: none"> • Dynamic: can be changed without restarting the partition. • For systems running AIX, version 5.2 or later, and Linux, the minimum amount of memory that can be allocated to a partition on a POWER5 system is 128 MB.

Partitioning the server

		<ul style="list-style-type: none"> • Memory can be assigned in increments of 16 MB, 32 MB, 64 MB, 128 MB, and 256 MB. • Increment value is determined by the logical memory block size set in the ASMI.
I/O	Dynamic: can be changed without restarting the partition.	Dynamic: can be changed without restarting the partition.
Configuration data	Service processor and HMC contain LPAR configuration data.	Service processor and HMC contain LPAR configuration data.
Console support	HMC	<ul style="list-style-type: none"> • HMC • Integrated Virtualization Manager
Main storage dumps	System level dump	<ul style="list-style-type: none"> • Service processor dump • Platform system dump
Updates also known as program temporary fix (PTF)	HMC Software Maintenance application	<ul style="list-style-type: none"> • HMC • Service partition
Micro-Partitioning	Micro-Partitioning is not supported.	Micro-Partitioning is supported on AIX 5.3
Virtual Ethernet client	Virtual Ethernet client is not supported.	<p>Virtual Ethernet client is supported on AIX 5.3.</p> <p>On systems managed by the Integrated Virtualization Manager, there is a maximum of four virtual LANs with up to two virtual Ethernet adapters per logical partition. (The management partition manages and belongs to all virtual LANs.)</p>
Virtual SCSI client	Virtual SCSI client is not supported.	Virtual SCSI client is supported on AIX 5.3.
Partition Load Manager	Partition Load Manager is not supported.	Partition Load Manager is supported.
POWER5	Not supported.	POWER5 is supported.
Operating system availability	Operating system availability is supported.	Operating system availability is not supported.

Parent topic: [Logical partition functional differences](#)

Linux logical partition functional differences between server hardware and previous hardware models

Use this information to see the new or significantly changed enhancements to Linux partitioning.

Linux logical partition functions have new and changed technical enhancements for new server hardware.

To identify Linux logical partition technical enhancements running on server hardware compared to previous hardware models, see the following table:

On serversserver hardware systems, Linux logical partitions support these enhancements.

Function	Linux on servers server hardware	Linux on server server hardware and Linux on server server hardware
LPAR user interface	<ul style="list-style-type: none"> • Hardware Management Console (HMC) required to partition the server. • HMC remote command • Integrated Virtualization Manager 	<ul style="list-style-type: none"> • Hardware Management Console (HMC) required to partition the server. • HMC remote command • Integrated Virtualization Manager
LPAR authority	<p>User roles created using the HMC:</p> <ul style="list-style-type: none"> • Super administrator • Operator • Viewer • Product engineer • Service representative 	<p>User roles created using the HMC:</p> <ul style="list-style-type: none"> • Super administrator • Operator • Viewer • Product engineer • Service representative
Maximum number of partitions	<p>The maximum number of partitions supported depends on the number of processors in the server model.</p> <p>The maximum number of logical partitions on Integrated Virtualization Manager-managed systems is ten times the number of processors.</p> <p>The maximum number of logical partitions on Virtual Partition Manager-managed systems is four Linux partitions and one partition.</p>	<p>The maximum number of logical partitions supported equals 10 times the number of processors in the server model.</p> <p>The maximum number of logical partitions on Integrated Virtualization Manager-managed systems is ten times the number of processors.</p>
Types of partitions	<ul style="list-style-type: none"> • AIX service partition • Partition profiles • Full system profiles • System profiles <p>Profiles are HMC-specific.</p>	<ul style="list-style-type: none"> • AIX or Linux service partition • Partition profiles • System profiles <p>Profiles are HMC-specific.</p> <p>On Integrated Virtualization Manager-managed systems, the Virtual I/O Server logical partition is called the "management partition."</p>

	<p>On Integrated Virtualization Manager-managed systems, the Virtual I/O Server logical partition is called the "management partition."</p>	
<p>Partition management</p>	<ul style="list-style-type: none"> • The Hardware Management Console (HMC) enablement utilities provide base infrastructure for communicating between the partition and the HMC. • The ppc64utils and lsvpd packages enable basic serviceability actions. • The Service Resource Manager, Platform Error Log Analysis, and I/O Error Log Analysis work with the Service Focal Point on the HMC to service hardware errors in a partition. • Inventory Scout reports system configuration information and facilitates update of system firmware. • Service Agent provides 	<ul style="list-style-type: none"> • The Hardware Management Console (HMC) enablement utilities provide base infrastructure for communicating between the partition and the HMC. • The ppc64utils and lsvpd packages enable basic serviceability actions. • The Service Resource Manager, Platform Error Log Analysis, and I/O Error Log Analysis work with the Service Focal Point on the HMC to service hardware errors in a partition. • Inventory Scout reports system configuration information and facilitates update of system firmware. • Service Agent provides additional capability to monitor a system and report problems. • The Dynamic LPAR Resource Manager (DRM) must be used with Linux distributions Red Hat Enterprise Linux version 4 or SUSE Linux Enterprise Server 9 to perform DLPAR processor and I/O capabilities. <p>See the Service and productivity tools Web site at https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html for more information.</p>

	<p>additional capability to monitor a system and report problems.</p> <p>See the Linux on POWER Service Diagnostic Web site for more details.</p>	
Processors	<ul style="list-style-type: none"> Processors on servers boxes cannot be changed dynamically. Processors cannot be shared among multiple partitions. 	<ul style="list-style-type: none"> Processors can be changed without restarting the partition with Linux distributions Red Hat Enterprise Linux version 4 and SUSE Linux Enterprise Server 9 and using the DLPAR Resource Manager (DRM). See the Service and productivity tools Web site at https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html for more information. Can be shared among multiple partitions. Sharing mode of capped and uncapped. Powered-off partition using dedicated processors can have its processors available to shared processing pool.
Memory	Assigned in 256 MB increments.	The xx value is set through Logical Memory Block Size in ASMI.
Virtual Ethernet	On Integrated Virtualization Manager-managed systems, there's a maximum of four virtual LANs, with up to two virtual Ethernet adapters per LPAR. (The management partition manages and belongs to all virtual LANs.)	Up to 4094 networks. On Integrated Virtualization Manager-managed systems, there's a maximum of four virtual LANs, with up to two virtual Ethernet adapters per LPAR. (The management partition manages and belongs to all virtual LANs.) On Virtual Partition Manager-managed systems, there's a maximum of four virtual Ethernet adapters per LPAR.
Virtual I/O		Supports virtual SCSI disk, CD, and tape attached to a dedicated Virtual I/O Server partition (VIOS) or i5OS partition on the same system.
DLPAR	Supports DLPAR processors and I/O when running Linux distributions Red Hat Enterprise Linux version 4 or SUSE Linux Enterprise Server 9 and the DLPAR Resource Manager (DRM).	Supports DLPAR processors and I/O when running Linux distributions Red Hat Enterprise Linux version 4 or SUSE Linux Enterprise Server 9 and the DLPAR Resource Manager (DRM). See the Service and productivity tools Web site at https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html for more information.
Micro-partitioning		Supports capped and uncapped sharing of all processors in the partition.
Fixes also known as program temporary fix (PTF)	Primary partition	

Types of logical partition configurations

This topic provides an overview of the types of logical partition configurations including the manufacturing default configuration, and the Virtual I/O Server partition.

- **Manufacturing default configuration**
The manufacturing default configuration is the initial partition setup of the managed system as received from your service provider.
- **Virtual I/O Server partition**
The Virtual I/O Server provides virtual SCSI and shared Ethernet capability to client logical partitions.

Parent topic: [Concepts for partitioning the server](#)

Manufacturing default configuration

The manufacturing default configuration is the initial partition setup of the managed system as received from your service provider.

When your system is in the manufacturing default configuration, you can install an operating system on the managed system and use the managed system as a nonpartitioned server. In this state, you do not need to manage the system using a Hardware Management Console (HMC).

If you choose to attach an HMC to the managed system for reasons other than partitioning (such as to activate Power On Demand), then the HMC displays the managed system as having one logical partition with one partition profile. All of the physical hardware resources on the system are automatically assigned to this logical partition, and any new physical hardware resources that are added to the managed system are added automatically to this logical partition. The name of the logical partition is the serial number of the managed system, and the name of the partition profile is `default`. You do not need to make any partitioning changes on the server if you do not want to do so.

However, if you use the HMC to create, delete, change, copy, or activate any logical partitions or partition profiles on the managed system, the system will then be in partition mode. You must then use the HMC to manage the managed system. If a managed system is managed using an HMC, and you want to return the managed system to a nonpartitioned state or if you want to partition the managed system with the Integrated Virtualization Manager, then you must follow a special procedure to reset the server. For more information about this procedure, see [Resetting the server to a nonpartitioned configuration](#).

Managed systems that are partitioned using the Integrated Virtualization Manager are not managed with an HMC. If a managed system is managed using the Integrated Virtualization Manager, then you do not need to reset the server to return the managed system to a nonpartitioned state. Also, you do not need to reset the server if you want to switch from using the Integrated Virtualization Manager to using an HMC. To switch to using an HMC, back up the data on each logical partition, attach the HMC to the server, create the logical partitions, and restore the data to the storage assigned to each logical partition.

Parent topic: [Types of logical partition configurations](#)

Virtual I/O Server partition

The Virtual I/O Server provides virtual SCSI and shared Ethernet capability to client logical partitions.

The Virtual I/O Server is installed in its own partition. It allows a physical adapter with attached disks on the Virtual I/O Server partition to be shared by one or more partitions, enabling client logical partitions to consolidate, and potentially minimize, the number of physical adapters required. It also facilitates the sharing of physical Ethernet adapters, allowing multiple client logical partitions to share a single Ethernet adapter.

For more information about the Virtual I/O Server, see [Using the Virtual I/O Server](#).

Parent topic: [Types of logical partition configurations](#)

Minimum hardware configuration requirements for logical partitions

This topic provides an overview of minimum hardware configuration requirements for AIX and Linux logical partitions.

- **Minimum configuration requirements for an AIX logical partition**
This topic lists the minimum configuration requirements for an AIX logical partition.
- **Minimum configuration requirements for a Linux partition**
Learn about configuration requirements for Linux on server hardware systems.

Parent topic: [Concepts for partitioning the server](#)

Minimum configuration requirements for an AIX logical partition

This topic lists the minimum configuration requirements for an AIX logical partition.

Every AIX logical partition requires the following hardware resources:

Minimum processor requirements

1 dedicated processor or 0.1 processing unit

Minimum memory requirements

128 MB

I/O requirements

- Physical or virtual storage adapter (SCSI card)
- Physical or virtual network adapter
- 2 GB storage

Parent topic: [Minimum hardware configuration requirements for logical partitions](#)

Minimum configuration requirements for a Linux partition

Learn about configuration requirements for Linux on server hardware systems.

Each Linux logical partition requires the following minimum hardware resources.

Minimum processor requirements

1 dedicated processor or 0.1 processing unit

Minimum memory requirements

128 MB

Minimum I/O requirements

- Storage adapter (physical or virtual)
- Network adapter (physical or virtual)
- Approximately 1 GB storage

Parent topic: [Minimum hardware configuration requirements for logical partitions](#)

Physical and virtual hardware resources

This topic provides an overview of the physical and virtual hardware resources used for AIX and Linux logical partitions including virtual adapters, processors, memory, expansion units, console options, I/O devices, and how each operating system implements virtual resources.

- **Virtual adapters**
On server hardware systems, virtual adapters, including Ethernet, Small Computer Systems Interface (SCSI), and serial, interact with the operating system like any other adapter, except that they are not physically present.
- **How each OS implements virtual resources**
This topic describes how AIX and Linux logical partitions use virtual resources to share system resources and to communicate with each other.
- **Processors**
This information describes how you can set the managed system to distribute its processing capacity among the logical partitions on the managed system.
- **Memory**
Processors use memory to temporarily hold information. Memory requirements for partitions depend on partition configuration, I/O resources assigned, and applications used.
- **Expansion unit**
You can add expansion units to many of the models to support additional features and devices. If you want to create logical partitions on your server, you may need to add an expansion unit that contains the additional hardware that you need for each logical partition.
- **Terminal options for logical partitions**
You can initiate a terminal session to the logical partitions on your managed system using a variety of methods. Your choice of terminal depends on your operating system and business needs.
- **I/O devices**
This topic discusses the different types of I/O devices used on your managed system, and how the I/O devices are added to logical partitions.

Parent topic: [Concepts for partitioning the server](#)

Virtual adapters

On server hardware systems, virtual adapters, including Ethernet, Small Computer Systems Interface (SCSI), and serial, interact with the operating system like any other adapter, except that they are not physically present.

The system administrator uses the `lscfg`, `lscfg -l`, or `lscfg -l -x` to create virtual adapters in order to use virtual I/O devices. Adapters can be added while the system is running. The virtual adapters are recorded in system inventory and management utilities. Converged location codes can be used to correlate operating-system level or partition-level software entities, such as `eth0`, `en0`, and `CMN21`, to adapters. Similarly, the Ethernet adapters are visible in the same way as physical Ethernet adapters.

By default, Virtual Ethernet Media Access Control (MAC) addresses are created from the locally administered range. Using the default MAC addresses, it is possible that different server hardware servers will have virtual Ethernet adapters with the same addresses. This can present a problem if multiple, virtual networks are bridged to the same physical network.

If a server partition providing I/O for a client partition fails, the client partition might continue to function, depending on the significance of the hardware it is using. For example, if one partition is providing the paging volume for another partition, a failure of the partition providing that particular resource will be significant to the other partition. However, if the shared resource is a tape drive, a failure of the server partition providing the resource will have only minimal effects on the client partition.

Virtual I/O client support

The following table summarizes operating system support for using virtual I/O devices.

	Virtual console	Virtual Ethernet	Virtual disk	Virtual CD	Virtual tape
AIX	Yes	Yes	Yes	Yes (when managed)	Not supported
Linux	Yes	Yes	Yes	Yes	Yes

AIX partitions support booting from virtual devices, including disk boot from virtual disk or network boot from virtual Ethernet.

The firmware running in AIX and Linux logical partitions recognizes virtual I/O and can start the partition from virtual I/O. IPL can be either from the network over virtual Ethernet, or from a device such as virtual disk or virtual CD.

Virtual I/O server support

The following table summarizes operating system support for providing virtual I/O to partitions.

	Virtual CD	Virtual console	Virtual disk	Virtual tape
Linux	Yes	Yes	No	No
Virtual I/O Server	Yes	Yes	Yes	Not supported

Version 1.1 of the Virtual I/O Server (08/2004) provides SCSI disk and shared Ethernet adapter function to client partitions.

Virtual I/O configuration is a combination of and operating system configuration, whereas adapter configurations are made in the . They include:

- Creation of virtual Ethernet adapters
- Creation of server adapters
- Creation of client adapters
- Determination of whether a partition can start and shut down another partition

After the configuration has been made, the partition is aware that virtual adapters and operating system configurations can be made. For Linux partitions, virtual adapters are listed in the device tree. The device tree contains adapters, not the devices under the adapter.

Parent topic: [Physical and virtual hardware resources](#)

How each OS implements virtual resources

This topic describes how AIX and Linux logical partitions use virtual resources to share system resources and to communicate with each other.

- **[How AIX implements virtual resources](#)**
AIX logical partitions can use virtual storage and networking devices that are made available by the Virtual I/O Server.
- **[How Linux implements virtual resources](#)**
Linux logical partitions can use virtual storage and networking devices that are made available by the Virtual I/O Server.

Parent topic: [Physical and virtual hardware resources](#)

How AIX implements virtual resources

AIX logical partitions can use virtual storage and networking devices that are made available by the Virtual I/O Server.

Devices and adapters are assigned to the Virtual I/O Server which can then make those devices available for multiple AIX logical partitions to share.

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

Parent topic: [How each OS implements virtual resources](#)

How Linux implements virtual resources

Linux logical partitions can use virtual storage and networking devices that are made available by the Virtual I/O Server.

Devices and adapters are assigned to the Virtual I/O Server which can then make those devices available for multiple Linux logical partitions to share.

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

Parent topic: [How each OS implements virtual resources](#)

Processors

This information describes how you can set the managed system to distribute its processing capacity among the logical partitions on the managed system.

A *processor* is a device that processes programmed instructions. The more processors that you assign to a logical partition, the greater the number of concurrent operations that the logical partition can run at any given time.

You can set up a logical partition to use either shared processors from a shared processor pool or processors that are dedicated to that logical partition. If a logical partition uses dedicated processors, then you must assign processors (in increments of whole numbers) to the logical partition. A logical partition that uses dedicated processors cannot use any processing capacity beyond the processors that are assigned to the logical partition, and no other logical partition can use the dedicated processors that are assigned to that logical partition.

All processors that are not dedicated to specific logical partitions are placed in the shared processor pool. The shared processor pool can be used by logical partitions that are set to use shared processors. You can set a logical partition that uses shared processors to use as little as 0.10 processing units, which is approximately one tenth of the processing capacity of a single processor. You can specify the number of processing units to be used by a shared processor logical partition down to the hundredth of a processing unit. Also, you can set a shared processor logical partition so that, if the logical partition requires more processing capacity than its assigned number of processing units, the logical partition can use unused processing units from the shared processor pool. (Some server models might require you to enter an activation code before you can create logical partitions that use shared processors.)

Unless the operating system or server model limits the number of processors that a logical partition can use, you can assign up to the entire processing capacity on the managed system to a single logical partition.

- One logical partition that uses one dedicated processor.

- One logical partition that uses uncapped shared processing units, with a maximum of one virtual processor for the logical partition.
- Up to ten logical partitions that use capped shared processing units, with a minimum of 0.10 processing units for each logical partition.
- One logical partition that uses one or two dedicated processors.
- Two logical partitions that use one dedicated processor each.
- One logical partition that uses uncapped shared processing units, with a maximum of two virtual processors for the logical partition.
- Two logical partitions that use uncapped shared processing units, with a maximum of one virtual processor for each logical partition.
- One logical partition that uses one dedicated processor and one logical partition that uses uncapped shared processing units, with a maximum of one virtual processor for the logical partition that uses uncapped shared processing units.
- One logical partition that uses one dedicated processor and from one to ten logical partitions that use capped shared processing units, with a minimum of 0.10 processing units for each logical partition that uses capped shared processing units.
- One logical partition that uses uncapped shared processors, with a maximum of one virtual processor for each logical partition that uses uncapped shared processing units, and from one to ten logical partitions that use capped shared processing units, with a minimum of 0.10 processing units for each logical partition that uses capped shared processing units.
- From one to twenty logical partitions that use capped shared processors, with a minimum of 0.10 processing units for each logical partition.

Automatic redistribution of work when a processor fails

If the server firmware detects that a processor is about to fail, or if a processor fails when the processor is not in use, then the server firmware creates a serviceable event. The server firmware can also deconfigure the failing processor automatically, depending upon the type of failure and the deconfiguration policies that you set up using the Advanced System Management Interface (ASMI). You can also deconfigure a failing processor manually using the ASMI. For more information, see [Setting deconfiguration policies](#) and [Deconfiguring hardware](#).

If a processor fails when the processor is in use, then the entire managed system shuts down. When a processor failure causes the entire managed system to shut down, the system deconfigures the processor and restarts. The managed system attempts to start the logical partitions that were running at the time of the processor failure with their minimum processor values. If the managed system does not have enough processor resources to start all of the logical partitions with their minimum processor values, then the managed system starts as many logical partitions as it can with their minimum processor values. If there are any processor resources remaining after the managed system has started the logical partitions, then the managed system distributes any remaining processor resources to the running logical partitions in proportion to their desired processor values.

- **Dedicated processors**
Dedicated processors are whole processors that are assigned to a single partition. This topic explains how to work with dedicated processors. This topic includes a flash demo that shows an example scenario of how a powered-off partition using dedicated processors will have its processors available to the shared processor pool.
- **Shared processors**
Shared processors are physical processors whose processing capacity is shared among multiple logical partitions. The ability to divide physical processors and share them among multiple logical partitions is known as *Micro-Partitioning*. This topic explains how to work with shared processors, the shared processor pool, capped partitions (partitions that cannot use more processor power than their assigned processing capacity), and uncapped partitions (partitions that can use more processor power than their assigned processing capacity). This topic includes flash demos that show example scenarios of how the system uses capped and uncapped partitions and the shared processing pool to optimize processing power.

Parent topic: [Physical and virtual hardware resources](#)

Dedicated processors

Dedicated processors are whole processors that are assigned to a single partition. This topic explains how to work with dedicated processors. This topic includes a flash demo that shows an example scenario of how a powered-off partition using dedicated processors will have its processors available to the shared processor pool.

If you choose to assign dedicated processors to a logical partition, you must assign at least one processor to that partition. Likewise, if you choose to remove processor resources from a dedicated partition, you must remove at least one processor from the partition.

Different partitioning tools assign dedicated processors to logical partitions in different ways. For more information on dedicated processor assignment by a specific partitioning tool, see [Partitioning tools](#).

By default, a powered-off logical partition using dedicated processors will have its processors available to the shared processor pool. When the processors are in the shared processor pool, an uncapped partition that needs more processing power can use the idle processing resources. However, when you power on the dedicated partition while the uncapped partition is using the processors, the activated partition will regain all of its processing resources. If you use the Hardware Management Console, you can prevent dedicated processors from being used in the shared processor pool by disabling this function in the partition properties panels.

Parent topic: [Processors](#)

Shared processors

Shared processors are physical processors whose processing capacity is shared among multiple logical partitions. The ability to divide physical processors and share them among multiple logical partitions is known as *Micro-Partitioning*. This topic explains how to work with shared processors, the shared processor pool, capped partitions (partitions that cannot use more processor power than their assigned processing capacity), and uncapped partitions (partitions that can use more processor power than their assigned processing capacity). This topic includes flash demos that show example scenarios of how the system uses capped and uncapped partitions and the shared processing pool to optimize processing power.

Note: For some models, Micro-Partitioning is an option for which you must obtain and enter an Advanced POWER Virtualization activation code. If you have not yet entered this activation code for these server models, see [Entering the activation code for Virtualization Engine technologies](#).

All physical processors that are not dedicated to specific logical partitions are grouped together in the *shared processor pool*. You can assign a specific amount of the processing capacity in the shared processor pool to each logical partition using the shared processor pool.

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

Some server models allow only a portion of the shared processor pool for use by logical partitions, so you are not always able to assign the full capacity of the shared processor pool to logical partitions. The LPAR Validation Tool (LVT) shows how much of the shared processor pool is available for partitioning use on each server model, so use the LVT to validate your partition plan. For more information on the LVT, see [LPAR Validation Tool](#).

Different partitioning tools assign shared processors to logical partitions in different ways. For more information on shared processor assignment by a specific partitioning tool, see [Partitioning tools](#).

Partitions in the shared processor pool can have a sharing mode of capped or uncapped. An *uncapped logical partition* is a logical partition that can use more processor power than its assigned processing capacity. The amount of processing capacity that an uncapped logical partition can use is limited only by the number of virtual processors assigned to the logical partition and the amount of unused processing capacity that is

available in the shared processor pool. In contrast, a *capped logical partition* is a logical partition that cannot use more processor power than its assigned processing capacity.

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

By default, logical partitions that use the shared processor pool are uncapped logical partitions. You can set a logical partition to be a capped logical partition if you do not want the logical partition to use more processing power than its assigned amount. For example, if you lease server capacity to other companies, you can create logical partitions that use the shared processor pool and set the logical partitions to capped. This prevents the logical partitions that belong to these companies from using more processing capacity than what is specified in the lease. Any unused processing resources will be used only by the uncapped logical partitions in the shared processor pool.

Although an uncapped logical partition can use more processor power than its assigned processing capacity, the uncapped logical partition can never use more processing units than its assigned number of virtual processors.

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

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

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

- **Virtual processors in the shared processor pool**

A *virtual processor* is a representation of a physical processor to the operating system of a logical partition that uses the shared processor pool. This topic explains how to set the minimum, desired, and maximum processing units and virtual processors for optimal partition performance.

Parent topic: [Processors](#)

Virtual processors in the shared processor pool

A *virtual processor* is a representation of a physical processor to the operating system of a logical partition that uses the shared processor pool. This topic explains how to set the minimum, desired, and maximum processing units and virtual processors for optimal partition performance.

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

operating system can easily calculate how many operations it can perform at a time by counting the whole number of processors that are available to it.

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

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

There are limits to the number of processing units that you can have for each virtual processor. The minimum number of processing units that you can have for each virtual processor depends on the server model. The maximum number of processing units that you can have for each virtual processor is always 1.00. This means that a logical partition cannot use more processing units than the number of virtual processors that it is assigned, even if the logical partition is uncapped.

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

Different partitioning tools assign virtual processors to logical partitions in different ways. For more information on virtual processor assignment by a specific partitioning tool, see [Partitioning tools](#).

Operating systems display virtual processors in system utilities and performance-monitoring programs in the same way that the operating systems would display physical processors. However, operating systems may differ in how they display processor information. For example, each physical POWER5 processor can run two threads simultaneously. When a logical partition uses the shared processor pool on a server that uses the POWER5 processor, each virtual processor can also run two threads simultaneously. System utilities in AIX take threads into account, so if an AIX system utility says that you have two processors, then you have two threads on one virtual processor. In contrast, system utilities in do not take threads into account, so if an system utility says that you have two processors, you have two virtual processors.

Parent topic: [Shared processors](#)

Memory

Processors use memory to temporarily hold information. Memory requirements for partitions depend on partition configuration, I/O resources assigned, and applications used.

Different partitioning tools assign memory to logical partitions in different ways. For more information on memory assignment by a specific partitioning tool, see [Partitioning tools](#).

Memory can be assigned in increments of 16 MB, 32 MB, 64 MB, 128 MB, and 256 MB. The default memory block size varies according to the amount of configurable memory in the system.

Amount of configurable memory	Default memory block size
Less than 4 GB	16 MB
Greater than 4 GB up to 8 GB	32 MB
Greater than 8 GB up to 16 GB	64 MB
Greater than 16 GB up to 32 GB	128 MB

Greater than 32 GB	256 MB
--------------------	--------

The memory block size can be changed by using the Logical Memory Block Size option in the Advanced System Management Interface (ASMI). The machine default value should only be changed under direction from your service provider. To change the memory block size, you must be a user with administrator authority, and you must shut down and restart the managed system for the change to take effect. If the minimum memory amount in any partition profile on the managed system is less than the new memory block size, you must also change the minimum memory amount in the partition profile.

Depending on the overall memory in your system and the maximum memory values you choose for each partition, the server firmware must have enough memory to perform logical partition tasks. Each partition has a Hardware Page Table (HPT). The size of the HPT is based on an HPT ratio of 1/64 and is determined by the maximum memory values you establish for each partition.

Server firmware requires memory to support the logical partitions on the server. The amount of memory required by the server firmware varies according to several factors. Factors influencing server firmware memory requirements include the following:

- Number of logical partitions
- Partition environments of the logical partitions
- Number of physical and virtual I/O devices used by the logical partitions
- Maximum memory values given to the logical partitions

Note: Firmware level updates can also change the server firmware memory requirements. Larger memory block sizes can exaggerate the memory requirement change.

Generally, you can estimate the amount of memory required by server firmware to be approximately 8% of the system installed memory. The actual amount required will generally be less than 8%. However, there are some server models that require an absolute minimum amount of memory for server firmware, regardless of the previously mentioned considerations.

When selecting the maximum memory values for each partition, consider the following:

- Maximum values affect the HPT size for each partition
- The logical memory map size for each partition

If the server firmware detects that a memory module is about to fail, or if a memory module fails when the memory module is not in use, then the server firmware creates a serviceable event. The server firmware can also deconfigure the failing memory module automatically, depending upon the type of failure and the deconfiguration policies that you set up using the Advanced System Management Interface (ASMI). You can also deconfigure a failing memory module manually using the ASMI.

If a memory module fails when the memory module is in use, then the entire managed system shuts down. When a memory module failure causes the entire managed system to shut down, deconfigure the failed memory module manually and restart the managed system. The managed system attempts to start the logical partitions that were running at the time of the memory module failure with their minimum memory values. If the managed system does not have enough memory to start all of the logical partitions with their minimum memory values, then the managed system starts as many logical partitions as it can with their minimum memory values. If there is any memory left over after the managed system is done starting logical partitions, then the managed system distributes any remaining memory resources to the running logical partitions in proportion to their desired memory values.

- **Memory requirements for an AIX logical partition**
This topic lists the memory requirements for an AIX logical partition.
- **Memory requirements for Linux logical partitions**
Linux logical partitions require a minimum of 128 MB of memory.

Parent topic: [Physical and virtual hardware resources](#)

Memory requirements for an AIX logical partition

This topic lists the memory requirements for an AIX logical partition.

The smallest partition that can be defined on a POWER5 system is 128 MB of memory. This is supported for partitions running AIX 5L Version 5.2 with the 5200-04 Technology Level or later, AIX 5.3, and Linux. There are implications on how big a partition can grow based on the amount of memory allocated initially. For partitions that are initially sized less than 256 MB, the maximum size is 16 times the initial size. For partitions initially sized 256 MB or larger, the maximum size is 64 times the initial size.

The smallest increment for adding or removing memory from a partition on POWER5 is 16 MB.

Parent topic: [Memory](#)

Memory requirements for Linux logical partitions

Linux logical partitions require a minimum of 128 MB of memory.

Parent topic: [Memory](#)

Expansion unit

You can add expansion units to many of the models to support additional features and devices. If you want to create logical partitions on your server, you may need to add an expansion unit that contains the additional hardware that you need for each logical partition.

Some expansion units can support only disk units (storage expansion unit), while others can support a variety of hardware (system expansion unit). Expansion units generally contain one or more system I/O buses with various I/O devices.

Parent topic: [Physical and virtual hardware resources](#)

Terminal options for logical partitions

You can initiate a terminal session to the logical partitions on your managed system using a variety of methods. Your choice of terminal depends on your operating system and business needs.

The following choices of terminal are available for AIX and Linux.

- Hardware Management Console (HMC)
 - ◆ Local connect at the HMC
 - ◆ Remote connect using Web-based System Manager
- Telnet
- OpenSSH with OpenSSL (included in the AIX expansion pack or the Linux distribution)
- Direct serial connection (ASCII terminal or PC connected with null modem cable)
- **Hardware Management Console terminal and console options**
The HMC provides virtual terminal emulation for AIX and Linux logical partitions
- **Integrated Virtualization Manager terminal and console options**
The *Integrated Virtualization Manager* is a browser-based system management interface for Virtual I/O Server that allows you to create and manage AIX and Linux logical partitions on a single server. The Integrated Virtualization Manager is supported only on select server models.

Parent topic: [Physical and virtual hardware resources](#)

Hardware Management Console terminal and console options

The HMC provides virtual terminal emulation for AIX and Linux logical partitions

The HMC connects to the server firmware. You use the HMC to specify to the server firmware how you want resources to be allocated among the logical partitions on the managed system. You also use the HMC to start and stop the logical partitions, update the server firmware code, manage Power On Demand, and transmit service information to service and support if there are any hardware problems with your managed system.

You can create virtual terminal sessions locally on the HMC by using the Server Management commands on the HMC. If you configure the HMC to allow remote access, you can also create virtual terminal sessions remotely through the HMC. You can create remote virtual terminal sessions on AIX and Linux logical partitions by using the Server Management commands on Web-based System Manager. In either case, you must configure the HMC to allow remote access, and you must configure encryption on the logical partitions for the session to be secure.

The HMC communicates with servers using service applications to detect, consolidate, and send information to your service provider for analysis.

For more information about adding and using the HMC, refer to [Managing the Hardware Management Console \(HMC\)](#).

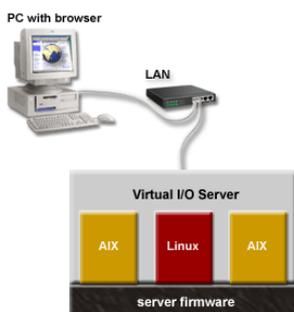
Parent topic: [Terminal options for logical partitions](#)

Integrated Virtualization Manager terminal and console options

The *Integrated Virtualization Manager* is a browser-based system management interface for Virtual I/O Server that allows you to create and manage AIX and Linux logical partitions on a single server. The Integrated Virtualization Manager is supported only on select 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 to Virtual I/O Server 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.



In this figure, you can see 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.

The Integrated Virtualization Manager does not have a command line interface that you can launch from the GUI. Instead, you can start a terminal session on the Virtual I/O Server logical partition from your PC. You can

start the terminal session using either *ssh* (if you have installed OpenSSL and Portable OpenSSH software on the management partition) or *telnet*. You can then use the *mkvt* command on the Virtual I/O Server command line interface. The *mkvt* command starts a virtual terminal session on any of the logical partitions on the managed system.

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

Parent topic: [Terminal options for logical partitions](#)

I/O devices

This topic discusses the different types of I/O devices used on your managed system, and how the I/O devices are added to logical partitions.

I/O devices allow your managed system to gather, store, and transmit data. I/O devices are found in the server unit itself and in expansion units and towers that are attached to the server. I/O devices can be embedded into the unit, or they can be installed into physical slots.

Not all types of I/O devices are supported for all operating systems or on all server models. For example, I/O processors (IOPs) are supported only on logical partitions. Also, Switch Network Interface (SNI) adapters are supported only on certain server models, and are not supported for logical partitions.

Parent topic: [Physical and virtual hardware resources](#)

Communications options for logical partitions

The types of communications options that you use depend on your business needs and the operating system you are running. You can establish interpartition communication using virtual Ethernet.

- [Virtual Ethernet](#)
This information provides an overview of virtual Ethernet and how AIX and Linux logical partitions use virtual Ethernet to establish interpartition communication.

Parent topic: [Concepts for partitioning the server](#)

Virtual Ethernet

This information provides an overview of virtual Ethernet and how AIX and Linux logical partitions use virtual Ethernet to establish interpartition communication.

- [Virtual Ethernet for AIX logical partitions](#)
- [Virtual Ethernet for Linux logical partitions](#)

Parent topic: [Communications options for logical partitions](#)

Virtual Ethernet for AIX logical partitions

Virtual Ethernet technology is supported on AIX 5.3 on POWER5 hardware. This technology enables IP-based communication between logical partitions on the same system using a VLAN capable software switch in POWER5 systems. Shared Ethernet Adapter technology (part of the optional Virtual I/O Server feature on POWER5 hardware) enables the logical partitions to communicate with other systems without assigning physical Ethernet slots to the logical partitions.

Virtual networking, along with other POWER5 virtualization technologies, offers greater flexibility in configuration scenarios. Workloads can be consolidated with more control over resource allocation. Network

availability might also be improved for more systems with fewer resources using a combination of virtual Ethernet, shared Ethernet and link aggregation in the Virtual I/O Server. When there are not enough physical slots to allocate a physical network adapter to each LPAR network, access using virtual Ethernet and a Virtual I/O Server is a preferable to IP forwarding because it does not complicate the IP network topology.

and servers support inter-LPAR communication using virtual networking. Virtual Ethernet adapters are connected to an IEEE 802.1q (VLAN)-style virtual Ethernet switch. Using this switch function, logical partitions can communicate with each other by using virtual Ethernet adapters and assigning VIDs (VLAN ID) that enable them to share a common logical network. The virtual Ethernet adapters are created and the VID assignments are done using the Hardware Management Console (HMC). The system transmits packets by copying the packet directly from the memory of the sender partition to the receive buffers of the receiver partition without any intermediate buffering of the packet.

The number of virtual Ethernet adapters for each LPAR varies by operating system. AIX 5.3 limits the number of virtual Ethernet adapters to 256 virtual Ethernet for each LPAR. Besides a PVID, the number of additional VID values that can be assigned for each virtual Ethernet adapter is 20, which indicates that each virtual Ethernet adapter can be used to access 21 networks. The HMC generates a locally administered Ethernet MAC address for the virtual Ethernet adapters so that these addresses do not conflict with physical Ethernet adapter MAC addresses. To ensure uniqueness among the virtual Ethernet adapters, the address generation is based on the system serial number, LPAR ID, and adapter ID.

When using the Integrated Virtualization Manager, only PVID is allowed (no additional VLANs), and only the PVID may only be 1-4. When using Virtual Partition Manager, each partition can have at most one Virtual Ethernet adapter for each PVID from 1-4.

For VLAN-unaware operating systems, each virtual Ethernet adapter is created with only a PVID (no additional VID values) and the POWER hypervisor ensures that packets have their VLAN tags removed before delivering to that LPAR. In the case of VLAN aware systems, you can assign additional VID values besides the PVID and the POWER hypervisor removes the tags of any packets that arrive with the PVID tag. Because the number of virtual Ethernet adapters supported for each LPAR is large, you can have multiple virtual Ethernet adapters with each adapter being used to access a single network and therefore assigning only PVID and avoiding the additional VID assignments. This also has the advantage that no additional VLAN configuration is required for the operating system using these virtual Ethernet adapters.

After a specific virtual Ethernet is enabled for a partition, a network device named `entX` is created in the partition. The user can then set up TCP/IP configuration similar to a physical Ethernet device to communicate with other partitions.

Parent topic: [Virtual Ethernet](#)

Virtual Ethernet for Linux logical partitions

Virtual Ethernet technology is supported on versions 2.4 and 2.6 of the Linux kernel on POWER5 hardware. Virtual Ethernet enables IP-based communication between logical partitions on the same system using a VLAN-capable software switch in POWER5 systems. The bridge module of the Linux kernel, along with the bridge-utils package, enables the logical partitions to communicate with other systems without assigning physical Ethernet slots to the logical partitions.

Virtual networking along with other POWER5 virtualization technologies offers greater flexibility in configuration scenarios. Workloads can be easily consolidated with more control over resource allocation. Network availability can also be improved for more systems with fewer resources using a combination of Virtual Ethernet, the bridge kernel module, and the bonding kernel module. When there are not enough physical slots to allocate a physical network adapter to each LPAR, network access using Virtual Ethernet and the bridge kernel module is preferable to IP forwarding because it does not complicate the IP network topology.

Server hardware supports inter-LPAR communication using virtual networking. Virtual Ethernet adapters are connected to an IEEE 802.1q (VLAN)-style virtual Ethernet switch. Using this switch function, logical partitions can communicate with each other by using virtual Ethernet adapters and assigning VLAN IDs (VID) that enable them to share a common logical network. The Virtual Ethernet adapters are created and the VID assignments are done using the Hardware Management Console. 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.

The number of Virtual Ethernet adapters for each LPAR varies by operating system. Version 2.4 of the Linux kernel supports up to 100 Virtual Ethernet adapters while version 2.6 of the Linux kernel can support up to 32 768 Virtual Ethernet adapters. Besides a Primary VID (PVID), the number of additional VID values that can be assigned for each Virtual Ethernet adapter is 19, which indicates that each Virtual Ethernet adapter can be used to access 20 networks. The HMC generates a locally administered Ethernet MAC address for the Virtual Ethernet adapters so that these addresses do not conflict with physical Ethernet adapter MAC addresses. To ensure uniqueness among the Virtual Ethernet adapters, the address generation is based on the system serial number, LPAR ID, and adapter ID.

When using the Integrated Virtualization Manager, only PVID is allowed (no additional VLANs), and only the PVID may only be 1-4. When using Virtual Partition Manager, each partition can have at most one Virtual Ethernet adapter for each PVID from 1-4.

For VLAN-unaware operating systems, each Virtual Ethernet adapter is created with only a PVID (no additional VID values) and the POWER hypervisor will ensure that packets have their VLAN tags removed before delivering to that LPAR. In the case of VLAN aware systems, such as Linux with the `vlan` module, you can assign additional VID values besides the PVID and the POWER hypervisor removes the tags of any packets that arrive with the PVID tag. Because the number of Virtual Ethernet adapters supported for each LPAR is large, you can have multiple Virtual Ethernet adapters with each adapter being used to access a single network and therefore assigning only PVID and avoiding the additional VID assignments. This also has the advantage that no additional VLAN configuration is required for the operating system using these Virtual Ethernet adapters.

After a specific virtual Ethernet is enabled for a partition, a network device named `ethX` is created in the partition. The user can then set up TCP/IP configuration similar to a physical Ethernet device to communicate with other partitions.

Parent topic: [Virtual Ethernet](#)

Logical partition utilities

This topic provides an overview of the utilities available to manage logical partitions including the Partition Load Manager for AIX.

- **Partition Load Manager for AIX**
The Partition Load Manager provides processor and memory resource management and monitoring across logical partitions (LPARs) within a single central processor complex (CPC).

Parent topic: [Concepts for partitioning the server](#)

Partition Load Manager for AIX

The Partition Load Manager provides processor and memory resource management and monitoring across logical partitions (LPARs) within a single central processor complex (CPC).

Partition Load Manager allows you to more effectively use resources by allowing you to set thresholds for designated resources. When a threshold is exceeded, Partition Load Manager can try to assign resources to that LPAR by using resources assigned to other LPARs that are not being used.

Partition Load Manager is available as part of the Advanced POWER Virtualization feature.

When the Partition Load Manager resource manager starts, it registers several events on every required LPAR node. The following events are registered individually on all of the managed LPAR nodes:

- Memory-page-steal high thresholds
- Memory-usage high thresholds and low thresholds
- Processor-load-average high thresholds and low thresholds

The Partition Load Manager resource manager tracks these threshold values. Every time a threshold is exceeded, Partition Load Manager receives a Resource Monitoring and Control (RMC) event. When a high threshold is exceeded, the node needs more resources. Alternately, when a low threshold is crossed, the

node has more resources available than it is currently using.

When a node requests additional resources, Partition Load Manager determines whether the node can accept additional resources. If the node can accept additional resources, Partition Load Manager conducts a search for available resources. Such additional resources can be found in the following places:

- The *free pool*, the list of currently unused resources maintained by Partition Load Manager. These resources are reallocated freely.
- Nodes that have indicated through events that they can release resources. These resources are removed from the node that does not require them and reallocated to the node requesting additional resources.
- Taken away from a node that has a lesser need for the resource, or a lower priority, than the node requesting the resource. These resources are removed from the node that has lower priority and reallocated to the node requesting additional resources.

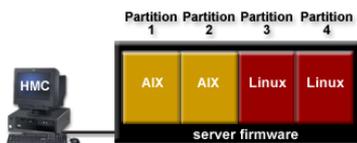
Determining which node is more or less deserving of resources is primarily done by taking into account certain values defined in what is known as a *policy file*. This policy file details partitions, their entitlements, their thresholds, and organizes the partitions into groups. Every node, but not every LPAR, managed by Partition Load Manager must be defined in the policy file along with several associated attribute values. Some of the attributes that are associated with the node are the maximum, minimum, and guaranteed resource values, variable share values, and so on. These are the attributes taken into consideration by Partition Load Manager when a decision is made as to whether a resource is reallocated from one LPAR to another.

For example, a machine is likely to lose its resource to a node with a higher variable shares attribute value if that machine has a lesser variable shares attribute value and currently has more resource than the guaranteed resource value given in the policy file.

Parent topic: [Logical partition utilities](#)

Partitioning with an HMC

The *Hardware Management Console (HMC)* is a system that controls managed systems, including the management of logical partitions and use of Power On Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information for analysis.



In this figure, you can see the logical partitions and the server firmware on the server. The *server firmware* is code that is stored in flash memory on the server. The server firmware directly controls the resource allocations on the server and the communications between logical partitions on the server. The HMC connects with the server firmware and specifies how the server firmware allocates resources on the server.

If you use a single HMC to manage a server, and the HMC malfunctions or becomes disconnected from the server firmware, then the server continues to run, but you will not be able to change the logical partition configuration of the server. If desired, you can attach an additional HMC to act as a backup and to provide a redundant path between the server and service and support.

When you configure logical partitions using the HMC, you must create at least one partition profile for each logical partition. A *partition profile* is a record on the HMC that specifies a resource configuration and activation state for the logical partition. Each logical partition has a default partition profile. If desired, you can create additional partition profiles with different resource configurations. When you activate a logical partition using the HMC, you choose a partition profile, and the server firmware starts the logical partition according to the specifications contained in the selected partition profile. For more information on partition profiles, see [Partition profile](#).

To simplify the process of starting an entire system configuration, you can create system profiles. A *system profile* is a record on the HMC that contains an ordered list of partition profiles. When you activate a system profile from the HMC, the server firmware activates the partition profiles in the system profile in the order in which the partition profiles are listed. For more information on system profiles, see [System profile](#).

Partitioning using the HMC is supported on all models and servers server models, although some models require you to enter an Advanced POWER Virtualization Technologies enablement code before partitioning the server.

- **Creating partitions from a system plan**

A *system plan* is a specification of the hardware and the logical partitions contained in one or more systems. Each system plan is contained within a *system-plan file*. You can import a system-plan file into a Hardware Management Console (HMC) and deploy the system plan to managed systems that are managed by the HMC. When you deploy a system plan to a managed system, the HMC creates the logical partitions specified in the system plan on the managed system. Depending on how the logical partitions are related to each other in the system plan, you can create all the logical partitions specified within the system plan, or you can choose which logical partitions you want to create from the system plan.

- **Partitioning for AIX**

Partitioning for AIX includes information about configuring and managing partitions that run the AIX operating system on ESCALA PL Series hardware.

- **Partitioning for Linux**

An overview of the process related to Linux partitioning.

Parent topic: [Partitioning the server](#)

Creating partitions from a system plan

A *system plan* is a specification of the hardware and the logical partitions contained in one or more systems. Each system plan is contained within a *system-plan file*. You can import a system-plan file into a Hardware Management Console (HMC) and deploy the system plan to managed systems that are managed by the HMC. When you deploy a system plan to a managed system, the HMC creates the logical partitions specified in the system plan on the managed system. Depending on how the logical partitions are related to each other in the system plan, you can create all the logical partitions specified within the system plan, or you can choose which logical partitions you want to create from the system plan.

You can use the `mksysplan` command in the HMC command line interface to create a system plan based upon the logical partition configuration of one of the managed systems that are managed by the HMC. You can then use the system plan to create identical logical partition configurations on managed systems with identical hardware. For more information about how you can create a system plan based upon an existing logical partition configuration, see the man page for the `mksysplan` command.

System plans are designed to create logical partitions on new managed systems that do not already have logical partitions created on them. 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 that are not in the Machine Default Configuration. Also, if you deploy a system plan in stages, and you change the resource allocations on the logical partitions on the managed system between stages, you can avoid validation problems in later stages by re-creating the system plan with the changed resource allocations.

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 import the file into the HMC using either of the following methods:

- 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.

After you import the system-plan file into an HMC, you can deploy the system plan within that file to the managed systems that are managed by the HMC.

The system-plan file uses the file suffix of `.sysplan`.

When you deploy a system plan, the HMC validates the system plan. The HMC deploys a system plan to a managed system only if the system plan schema 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.

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
- 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 would cause the system to fail validation. A system plan would 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. Internal Tnn slots are not validated either.

When validating an existing logical partition, the HMC validates the following for that logical partition. Validation fails for the existing logical partition if any step fails. 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.

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 looks for 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
- Physical I/O slot assignments
- Virtual I/O adapters and settings

For example, 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 invalid. 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 invalid. 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.

Also, the HMC does not install the operating systems on the logical partitions. Because of this, the HMC is also unable to configure virtual I/O adapters within the operating systems so that logical partitions can provide virtual storage resources to other logical partitions. After the system plan is deployed, you must perform these tasks manually.

- **Importing a system-plan file**

Use this procedure to import a system-plan file into a Hardware Management Console (HMC). You can import system-plan files from media that is mounted on the HMC (such as optical discs or USB drives), or you can import system-plan files from a remote site using FTP. You must import a system-plan file into an HMC before you can deploy the system plan in the system-plan file to a managed system. System plan files must have the extension **.sysplan** to be imported.

- **Deploying a system plan**

Use this procedure to deploy all or part of a system plan to a managed system using a Hardware

Management Console (HMC). When you deploy a system plan, the HMC creates logical partitions on the managed system according to the specifications in the system plan.

- **Exporting a system-plan file**

Use this procedure to export a system-plan file that is stored on a Hardware Management Console (HMC). You can export system-plan files to media that is mounted to the HMC (such as optical discs or USB drives), or you can export system-plan files to a remote site using FTP. This allows you to import the system-plan file into a different HMC and deploy the system plan to a managed system with identical hardware.

Parent topic: [Partitioning with an HMC](#)

Importing a system-plan file

Use this procedure to import a system-plan file into a Hardware Management Console (HMC). You can import system-plan files from media that is mounted on the HMC (such as optical discs or USB drives), or you can import system-plan files from a remote site using FTP. You must import a system-plan file into an HMC before you can deploy the system plan in the system-plan file to a managed system. System plan files must have the extension **.sysplan** to be imported.

If you want to import the system-plan file from a remote site using FTP, the HMC must have a network connection to the remote site, an FTP server must be active on the remote site, and port 21 must be open on the remote site.

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 an HMC, complete the following:

1. If you want to import the system-plan file from media, mount the media on the HMC. Complete the following:
 - a. Insert the media into the HMC or connect the media to the HMC.
 - b. On your HMC desktop (outside of any of the displayed windows), right-click Terminal > rshterm. The Restricted shell command line interface displays.
 - c. Type the command `mount mountpoint`, where mountpoint is the location at which the media is located. You can typically use `mount /mnt/cdrom` to mount an optical drive or `mount /media/sda1` to mount a USB drive.
 - d. Type the command `lsmediadev` to list the mounted drives on the HMC and verify that the media is mounted and ready.
2. In the navigation area of your HMC, select System Plans.
3. In the contents area, select Import System Plan.
4. Enter the name of the system-plan file into the System plan file name field. The name of the system-plan file must end with the .sysplan file name suffix.
5. Specify whether you are importing the system-plan file from locally mounted media or from a remote FTP site.

From where are you importing the system-plan file?	Complete the following:
Locally mounted media	<ol style="list-style-type: none"> a. Select Import from media. b. Enter the path in which the system-plan file is located into the Directory field.
Remote FTP site	<ol style="list-style-type: none"> a. Select Import from a remote site. b. Enter the host name or IP address of the remote FTP site into the Remote site hostname field. c. Enter the user ID to use to access the remote FTP site into the User id field. d. Enter the password to use to access the remote FTP site into the Password field. e. Enter the path in which the system-plan file is located on the remote FTP site into the Remote directory field. If you do not

	enter a path, the HMC looks at the default path specified on the remote FTP site.
--	---

- Click Import. If the HMC returns an error, return to step 4 and verify that the information you entered on this window is correct.

When you are done, you can deploy the system plan in the system-plan file to a managed system that is managed by the HMC. For more information about how to deploy the system plan, see [Deploying a system plan](#). 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.

Parent topic: [Creating partitions from a system plan](#)

Deploying a system plan

Use this procedure to deploy all or part of a system plan to a managed system using a Hardware Management Console (HMC). When you deploy a system plan, the HMC creates logical partitions on the managed system according to the specifications in the system plan.

Before you deploy a system plan, complete the following:

- 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 more information about importing a system-plan file, see [Importing a system-plan file](#).
- Delete the logical partition that came with your server, and delete any other logical partition that is not in the system plan. Every server comes with one logical partition with one partition profile. The name of the logical partition is the serial number of the managed system, and the name of the partition profile is default. All of the physical hardware resources on the system are automatically assigned to this logical partition. The logical partition allows you to power on the server and verify that the physical hardware is connected and reporting to the server. After you are done verifying the physical hardware, you must delete this logical partition before the HMC allows you to deploy your system plan. Verify that this partition is the service partition for the managed system. It can not be deleted until it is no longer the service partition. Use the Server Management for the managed system and set the properties for the service partition field to Unassigned. Then remove the logical partition.

To remove the default partition profile, complete the following steps:

- In either the Web-based Systems Management or native HMC interface select Server and Partition > Server Management.
 - Right click on the managed system in the content pane, and select Properties.
 - Select Unassigned from the Service Partition box on the Properties panel, and click Ok.
 - On the main content panel, click the Plus icon for the managed system to expand the managed system and list the available partitions.
 - Right-click the default partition and select Delete, and then click Ok.
- Shut down any logical partitions that you have already deployed to the managed system from the system plan.
 - Ensure that you are not performing any other operations on the managed system using this HMC or any other HMC that is attached to the managed system.

To deploy a system plan, you must be a super administrator. For more information about user roles, refer to [Tasks and roles](#).

To deploy a system plan on a managed system using the HMC, complete the following:

- In the navigation area of your HMC, select System Plans.
- In the contents area, select Manage System Plans.

3. Select the system-plan file that contains the system plan that you want to deploy and click Deploy. If you are not certain which system-plan file to choose, you can select a system-plan file and click View to list the contents of a system-plan file in a browser window.
4. Verify that the system-plan file that you want to deploy is displayed in and click Next.
5. If there are multiple system plans within the system-plan file, the wizard prompts you to select the system plan that you want to deploy from that system-plan file. Select the system plan that you want to deploy in the System plan to deploy field and click Next.
6. Choose the managed system to which you want to deploy the system plan in Managed system 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 dialog that informs you of this. Click OK to continue or Cancel to select a different system plan.
7. Wait for the wizard to validate the managed system and its hardware against the system plan. The validation process can take several minutes.
8. If the validation process completes successfully, click Next. If the validation process does not complete successfully, correct the issues indicated by the dialog, click Cancel to exit the wizard, and restart this procedure from the beginning. If the validation failed, you should also run the following command on the invalid system: **mksysplan -m name_of_managed_system -f name_of_new_system_plan.sysplan.**

This creates a new system plan that can be viewed and compared to the and the old system plan to help diagnose any problems.
9. 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.
10. Click Next to continue.
11. Review the system deployment step order and click Finish. The HMC creates the specified logical partitions. This process can take several minutes.

After you have deployed the system plan, complete the following:

- Locate the physical disk I/O adapters that belong to each logical partition and verify that the disk drives that are attached to these physical I/O adapters will support your desired configuration for each logical partition.
- Install operating systems and software on the logical partitions. For installation procedures for the AIX, , and Linux operating systems, refer to [Installing operating systems](#).
- 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.

You can also remove the system plan from the HMC by selecting the system plan in the Manage System Plans window, clicking Remove, and clicking Remove System Plan to confirm.

Parent topic: [Creating partitions from a system plan](#)

Exporting a system-plan file

Use this procedure to export a system-plan file that is stored on a Hardware Management Console (HMC). You can export system-plan files to media that is mounted to the HMC (such as optical discs or USB drives), or you can export system-plan files to a remote site using FTP. This allows you to import the system-plan file into a different HMC and deploy the system plan to a managed system with identical hardware.

If you want to export the system-plan file to a remote site using FTP, the HMC must have a network connection to the remote site, an FTP server must be active on the remote site, and port 21 must be open on the remote site.

To export a system-plan file, you must be a super administrator. For more information about user roles, refer to [Tasks and roles](#).

To export a system-plan file that is stored on an HMC, complete the following:

1. If you want to export the system-plan file to media, mount the media on the HMC. Complete the following:
 - a. Insert the media into the HMC or connect the media to the HMC.
 - b. On your HMC desktop (outside of any of the displayed windows), right-click Terminal > rshterm. The Restricted shell command line interface displays.
 - c. Type the command `mount mountpoint`, where mountpoint is the location at which the media is located. You can typically use `mount /mnt/cdrom` to mount an optical drive or `mount /media/sda1` to mount a USB drive.
 - d. Type the command `lsmediadev` to list the mounted drives on the HMC and verify that the media is mounted and ready.
2. In the navigation area of your HMC, select System Plans.
3. In the contents area, select Manage System Plans.
4. Select the system-plan file that you want to export and click Export.
5. Specify whether you are exporting the system-plan file to locally mounted media or to a remote FTP site.

To where are you exporting the system-plan file?	Complete the following:
Locally mounted media	<ol style="list-style-type: none"> a. Select Export to media. b. Enter the path to which you want to export the system-plan file into the Directory field.
Remote FTP site	<ol style="list-style-type: none"> a. Select Export to a remote site. b. Enter the host name or IP address of the remote FTP site into the Remote site hostname field. c. Enter the user ID to use to access the remote FTP site into the User id field. d. Enter the password to use to access the remote FTP site into the Password field. e. Enter the path to which you want to export the system-plan file into the Remote directory field. If you do not enter a path, the HMC tries to export the system-plan file to the default path specified on the remote FTP site.

6. Click Export. If the HMC returns an error, verify that the information you entered on this window is correct. If necessary, click Cancel, return to step 4, and redo the procedure, ensuring that the information you specify at each step is correct.

If you exported the system-plan file to media, you can unmount the media by using the `umount` command in the HMC command line interface. You can then import the system-plan file into a different HMC so that you can deploy the system plan to managed systems that are managed by the other HMC. For more information about importing the system-plan file into a different HMC, see [Importing a system-plan file](#).

Parent topic: [Creating partitions from a system plan](#)

Related information for logical partitions

There are Web sites, and ESCALA Power5 Hardware Information topics that relate to the Partitioning the server topic.

For more information about AIX and Linux logical partitions, see:

- [Related information for AIX logical partitions](#)
This topic provides links to information related to AIX logical partitions.
- [Related information for Linux logical partitions](#)
Use this information to find related material for Linux logical partitions, including Web sites, , and ESCALA Power5 Hardware Information topics.

- [Related information for the Integrated Virtualization Manager](#)
Listed here are ESCALA Power5 Hardware Information topics that relate to the Integrated Virtualization Manager.
- [Related information for the Virtual I/O Server](#)
Find other information related to the Virtual I/O Server.

Parent topic: [Partitioning the server](#)

Related information for AIX logical partitions

This topic provides links to information related to AIX logical partitions.

Listed below are information center topics that relate to the Partitioning for AIX with an HMC topic. You can view or print any of the PDF files.

Other information

- [Logical partition functional differences between new hardware and previous hardware models](#)
- [Upgrades](#)
- [Partitioning for Linux with an HMC](#)
- [Managing your server](#)
- [Working with Power On Demand \(POD\)](#)
- [Overview of service and support](#)

Parent topic: [Related information for logical partitions](#)

Related information for Linux logical partitions

Use this information to find related material for Linux logical partitions, including Web sites, , and information center topics.

Redbooks

- [AIX and Linux Interoperability, SG24 6622 00](#) (4.3 MB)

ESCALA Power5 Hardware Information topics

- [Installing Linux](#)
- [Using the Virtual I/O Server](#)
- [Partitioning for AIX](#)
- [Managing your server](#)
- [Working with Power On Demand](#)
- [Customer service and support](#)

Saving PDF files

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

1. Right-click the PDF in your browser (right-click the link above).
2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you would like to save the PDF.
4. Click Save.

Downloading Adobe Reader

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

Parent topic: [Related information for logical partitions](#)

Related information for the Integrated Virtualization Manager

Listed here are ESCALA Power5 Hardware Information topics that relate to the Integrated Virtualization Manager.

Other information

- [Managing the Integrated Virtualization Manager](#)
- [Partitioning the server](#)
- [Creating a virtual computing environment](#)
- [Using the Virtual I/O Server](#)
- [Virtual I/O Server command descriptions](#)

Parent topic: [Related information for logical partitions](#)

Related information for the Virtual I/O Server

Find other information related to the Virtual I/O Server.

The following Web sites and ESCALA Power5 Hardware Information topics relate to the Virtual I/O Server topic.

Other information

- [Managing your server](#)
- [Customer service and support](#)

Saving PDF files

To save a PDF on your workstation for viewing or printing, use the following task:

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

Downloading Adobe Reader

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

Parent topic: [Related information for logical partitions](#)

Technical publication remarks form

Title : ESCALA POWER5 Hardware Information Partitioning the server

Reference N° : 86 A1 18EW 00

Date: July 2006

ERRORS IN PUBLICATION

--

SUGGESTIONS FOR IMPROVEMENT TO PUBLICATION

--

Your comments will be promptly investigated by qualified technical personnel and action will be taken as required.
If you require a written reply, please include your complete mailing address below.

NAME : _____ Date : _____

COMPANY : _____

ADDRESS : _____

Please give this technical publication remarks form to your BULL representative or mail to:

Bull - Documentation Dept.
1 Rue de Provence
BP 208
38432 ECHIROLLES CEDEX
FRANCE
info@frec.bull.fr

Technical publications ordering form

To order additional publications, please fill in a copy of this form and send it via mail to:

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

Phone: +33 (0) 2 41 73 72 66
FAX: +33 (0) 2 41 73 70 66
E-Mail: srv.Duplicopy@bull.net

CEDOC Reference #	Designation	Qty
-- -- []		
-- -- []		
-- -- []		
-- -- []		
-- -- []		
-- -- []		
-- -- []		
-- -- []		
-- -- []		
-- -- []		
-- -- []		
-- -- []		

[] : The latest revision will be provided if no revision number is given.

NAME: _____ Date: _____

COMPANY: _____

ADDRESS: _____

PHONE: _____ FAX: _____

E-MAIL: _____

For Bull Subsidiaries:

Identification: _____

For Bull Affiliated Customers:

Customer Code: _____

For Bull Internal Customers:

Budgetary Section: _____

For Others: Please ask your Bull representative.

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

REFERENCE
86 A1 18EW 00